arXiv:1603.07150v1 [cs.DL] 23 Mar 2016

The Anatomy of a Search and Mining System for Digital Archives

Martyn Harris, Mark Levene, Dell Zhang
Department of Computer Science and Information Systems
Birkbeck, University of London
London WC1E 7HX, UK
martyn, mark, dell @dcs.bbk.ac.uk

Dan Levene
History, School of Humanities
University of Southampton
Southampton SO17 1BF, UK
d.levene@soton.ac.uk

Abstract

Samtla (Search And Mining Tools with Linguistic Analysis) is a digital humanities system designed in collaboration with historians and linguists to assist them with their research work in quantifying the content of any textual corpora through approximate phrase search and document comparison. The retrieval engine uses a character-based n-gram language model rather than the conventional word-based one so as to achieve great flexibility in language agnostic query processing.

The index is implemented as a space-optimised character-based suffix tree with an accompanying database of document content and metadata. A number of text mining tools are integrated into the system to allow researchers to discover textual patterns, perform comparative analysis, and find out what is currently popular in the research community.

Herein we describe the system architecture, user interface, models and algorithms, and data storage of the Samtla system. We also present several case studies of its usage in practice together with an evaluation of the systems' ranking performance through crowdsourcing.

Keywords: Digital Humanities, Statistical Language Model, Information Retrieval, Text Analysis

1 Introduction

Digital textual representations of original historic documents are being made available thanks to the work of digital archiving projects [57]. As the wealth of digitised textual data increases it becomes more important to provide the appropriate tools for analysing and interrogating the sources. Humanities researchers are discovering how digital tools can become part of their methodology - even collating their source materials into a simple repository can help to speed up the access to resources which were otherwise stored in physical libraries. However, there are still barriers to adoption, including usability and the scope of the provided tools [35, 64]. In addition, there is a tendency to develop systems which are language specific and rely on part of speech taggers to identify all instances of a word, regardless of its morphology, in order to provide accurate recall. Furthermore, such systems may be tied to a particular document collection, for example, the works of William Shakespeare. As a result, when developing any tool for the humanities important aspects to consider are how to index, search, compare, and apply data mining tools to domain-specific corpora represented by a collection of documents grouped together for a specific research agenda.

Samtla has been designed to provide a research environment that is agnostic to the document collection and can therefore be used by a wide range of research groups whose work involves analysing digital representations of original source texts. It currently supports search, browsing, and analysis of texts through approximate phrase searches, related query and document recommendations, a document comparison tool, and community features in the form of popular queries and documents. Samtla's interface adopts the flat design principle, which reflects current trends [4, 5] in user interface design in terms of user interaction and the layout of components such as tool bars and informational side panels to promote familiarity with respect to applications the user may use regularly (including browsers, music and video players, cloud storage), and legibility in terms of centralising the content by presenting it clearly to users.

Samtla was developed in collaboration with historians and linguists to cater for their research needs in quantifying the content of textual corpora. We have adopted a Statistical Language Model (SLM) for information retrieval [67], and incorporated text mining tools into the system to allow researchers to go beyond a pure search and browse paradigm. Such an extended "search and research" model supports the discovery of patterns (known by historians as "formulae" or "parallel passages") that have significance in terms of their research goals. The "formulae" are reflected by textual fragments represented, for example, by set phrases or quotations. The main textual fragment is duplicated across several documents with slight variations resulting from differences in authorship, language change due to locality and time, which can manifest themselves through dialectal differences. The system has been designed to be applicable to any type of corpus in any language with little pre-processing, and to provide transparency in terms of its functionality, in order to help researchers adopt it as an integral

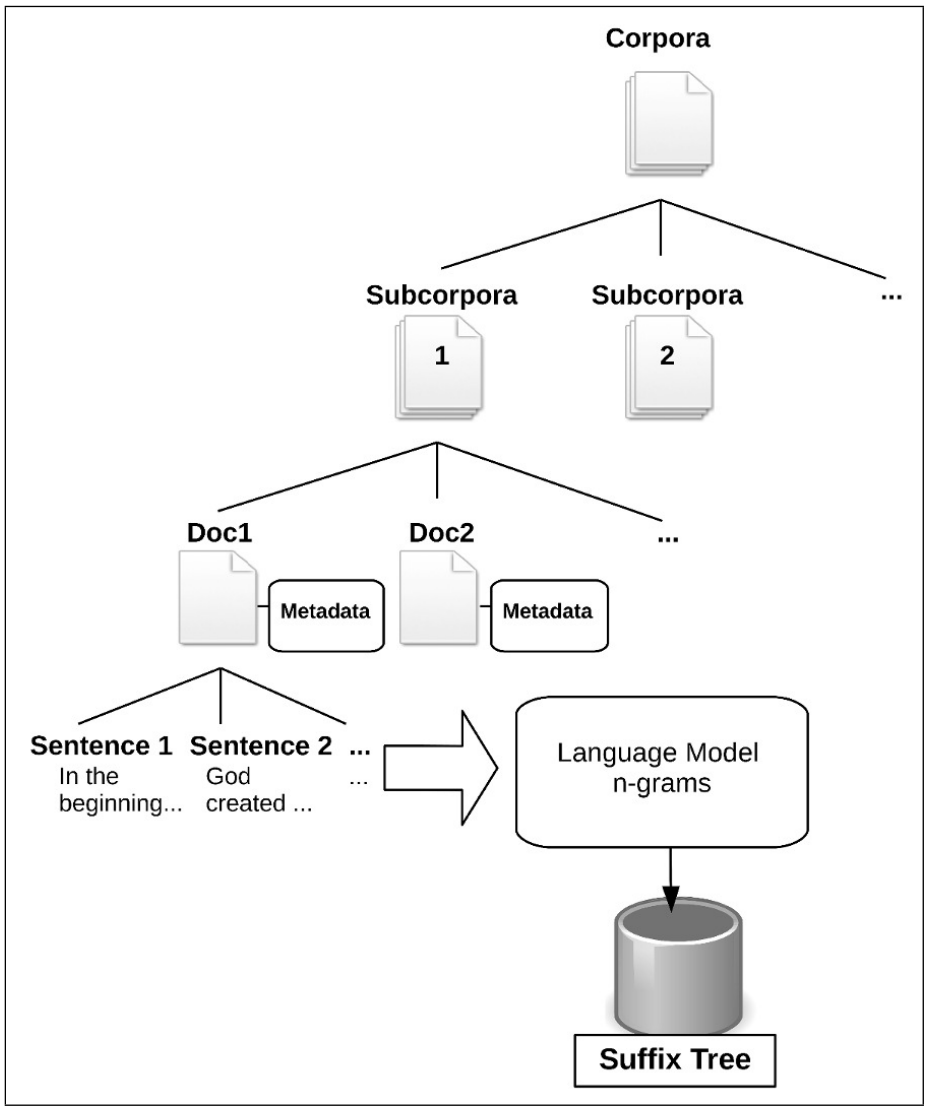

Figure 1: The corpus typology.

part of their research strategy [35, 64]. For this purpose, our n-gram statistical language model is character-based rather than word-based, which makes the query processing and further analysis flexible. For instance, consider the difficulty involved in indexing words from a collection of English news articles and those in the Chinese language; the latter of which has no whitespace equivalent word boundary marker as that present in English.

Figure 1 illustrates the basic structure of a corpus and the documents contained within. Some corpora maybe partitioned in to subcorpora, and documents may also come with metadata, either in terms of document features (title, pages, and length) or additional metadata items generated by the user group.

Samtla currently operates with five case study textual corpora: a collection of Aramaic texts from late antiquity, Italian and English translations of the writings of Giorgio Vasari, the Microsoft corpus of 68,000 scanned books, the King James Bible in English, and a test corpus of scanned Newspaper articles from the Financial Times as part of a pilot study in collaboration with the British Library. Samtla was developed to enable faster access to the document collection and to compliment existing methods adopted by our users for comparison and discovery of related documents and parallel text fragments.

In the following sections we discuss prominent tools that are currently available for researchers (Section 2), outline each of the system components described briefly above in more detail including the Samtla system architecture (Section 3), the statistical models underlying document scoring and recommendation, and a discussion of the chosen implementation methods used to structure and organise the data used by the system (Section 4). We introduce the tools that have been implemented to help researchers browse, view documents and related metadata, and compare shared-sequences between document pairs (Section 5). Next, we introduce the recommendation component of the system, which leverages data collected from the user query submissions and document views to generate recommended queries and documents to help users locate interests aspects of the collection (Section 5.5). Section 6 provides a description of the User Interface (*UI*) and discusses how the various tools have been incorporated in to the interface and how we expect users to navigate through the system. We also present case studies describing how the system is currently being used by researchers (Section 7), before moving to the results of a formal evaluation using a crowdsourcing platform (Section 8). Lastly, we conclude the paper with a summary of the work and future development plans (Section 9).

2 Related Work

There are a number of existing systems that provide state of the art tools in the Humanities. We discuss some of the systems that share common functionality, source material, or user groups with reference to Samtla.

The Bar-Ilan *Responsa* project, established in 1963, is one of the earliest examples of Humanities researchers adopting the use of computer-based methods for search, comparison, and analysis of Hebrew texts. The corpus spans approximately three thousand years, and includes the Mishnah, Talmud, Torah, and the Bible in Aramaic [27]. The Responsa environment is packaged on a CD-ROM and provides browsing and searching the corpus using keyword and

phrase search, comparison of parallel passages, author biographies, and user comment and annotation tools [16]. The Bar-Ilan *Responsa* project has made a considerable contribution to Computational Linguistics and information retrieval for Hebrew.

The 1641 depositions project at Trinity College Library (Dublin) [3] adopted IBM's *LanguageWare* [2] for the analysis of 31 volumes of books containing 19,010 pages of witness accounts reporting theft, vandalism, murder, and land taking during the conflicts between Catholics and Protestants in 17th century Ireland. *LanguageWare* provides text analysis tools for mining facts from large repositories of unstructured text. Its main features are dictionary and fuzzy look-up, lexical analysis, language identification, spelling correction, part-of-speech disambiguation, syntactic parsing, semantic analysis, and entity and relationship extraction. *LanguageWare* was chosen due to the complexity of the language contained in the documents, which have many spelling mistakes making analysis a complex task. *CULTivating Understanding Through Research and Adaptivity* (*CULTURA*) is a project related to the 1641 depositions, launched in 2011 [1]. *CULTURA* provides users with tools for normalising texts containing inconsistent spelling, entity and relationship extraction present within unstructured text, and social network analysis tools for displaying the entities and relationships from metadata through an interactive user environment.

Aside from tools developed as part of funded projects, we also see new applications for mobile and touch devices, developed specifically for exploring well known texts such as the Bible. One such tool is *Accordance* for *Apple iOS*[8], which is a Bible study application featuring exact and flexible query search tools, a browsing, a timeline for viewing when people lived and important events that took place, and an atlas view for exploring journeys and battles.

The systems descibed above are successful at providing the appropriate tools for specific groups of researchers to explore a particular collection of texts. The main issue, however, is the generalisability of their systems to other text domains and natural languages. These systems use language specific tools that operate at the word or morphological level requiring affix removal, tokenisation, lemmatisation, and part-of-speech tagging for normalising the texts and capturing all instances of a word in order to generate an accurate retrieval model. The question is whether they could be extended to languages with no clear word-level boundary markers (e.g. Chinese languages such as Mandarin), or without the complex rules necessary to identify affixes (e.g. Hebrew, Aramaic, Arabic, Italian, and Russian). Regardless of the approach adopted for text normalisation, these systems are by their nature language dependent. If we are to keep up with the volume of output generated by digitisation projects, then a new approach is necessary. The Samtla system was designed to address the need for tools for the Digital Humantities through the creation of a flexible language-independent framework for searching, browsing, and comparing documents in a text collection that could be generalised to any document collection

due to its data-driven design. Samtla shares many of the features and tools available in the systems outlined above, however it differs in many respects due to the language-independent framework that can be extended in many novel ways without changes to the underlying system each time a new corpus is introduced. This enables a Samtla system to be deployed relatively quickly, allowing document collections and archives to be unlocked to the general public, or for research once the digitised materials are made available. The philosophy underlying the development of Samtla has been to provide the basic tools first (browse, search, and comparison), and then to develop further features through consultation with our users to discover what tools they actually need in order to be able to carry out their research.

3 System Architecture

The Samtla system is a web-based application built on a client-server architecture, providing a platform-independent solution for its deployment through a web browser. The Samtla system operates with a single code-base, with the only corpus-dependent component being a wrapper function, which is responsible for parsing the documents or metadata to the system. This enables the system to be data-driven and allows upgrades or changes to the functionality of Samtla to be rolled-out simultaneously to each user group.

The client is represented by a web-based User Interface (UI), which sends requests to the server and renders the results within the browser. A central server stores all the data associated with the system, processes requests sent from the client, and responds with the appropriate data, for instance, a list of search query results.

Samtla can be viewed as *a Model-View-Controller* (MVC) design pattern [45], allowing the separation of the system by function. The advantage of adopting a MVC implementation is that changes to the UI are independent of the underlying logic of the system. Therefore, due to the separation of components, introducing new features and changes to the look or functionality of the UI can be easily implemented without affecting other components. An overview of Samtla is shown in Figure 2, where arrows in the diagram represent the flow of communication between the various components of the system.

The client-side of the system is represented by the *view* component (i.e. the UI) providing a web-based browser interface, which allows the user to interact with the system; see Section 6 for more detail on the UI. Such interactions cause events to be triggered and picked up by the *controller*, representing the program logic. Technically, the client communicates with the server through the controller using URL requests, which are mapped to an appropriate function call in the *model*. This instigates a change to the data or the retrieval of infor-

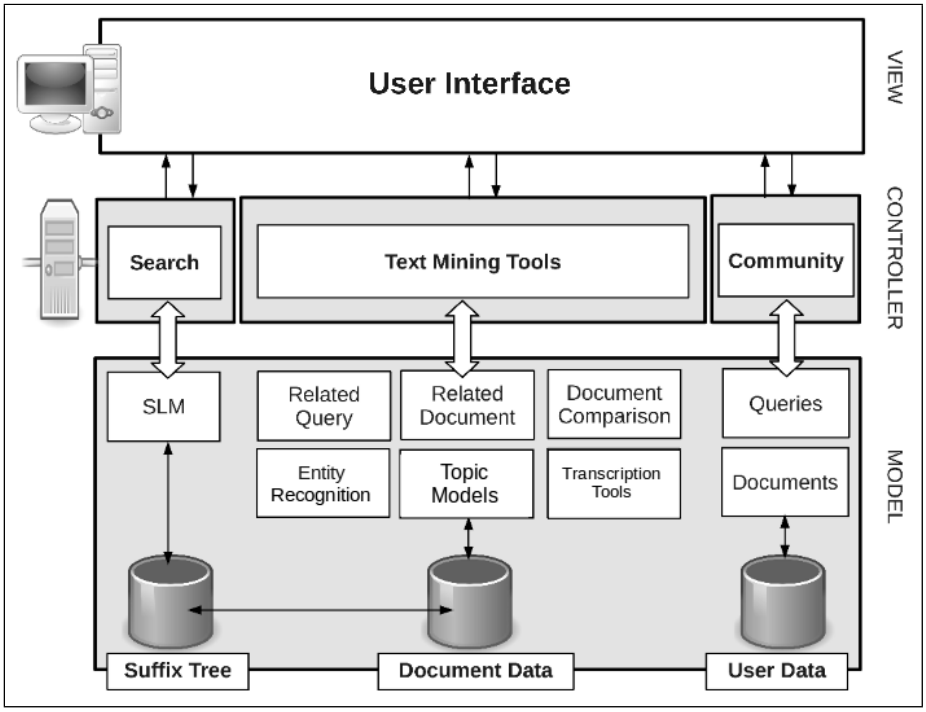

Figure 2: The Samtla architecture.

mation such as search query results or metadata. The *model* returns data to the controller to process, which passes the results to the view for rendering to the UI. The controller is implemented in the Django web framework [9], which processes client HTTP requests and sends a response. The data is passed between components using the JavaScript Object Notation (JSON) [10] format, which allows us to store data objects that can be further processed in the browser (i.e. for dynamic rendering of HTML snippets for the search results), or static HTML fragments which are rendered directly (i.e. the raw documents). The UI is developed in Javascript with JQuery [12], providing cross-browser support for the interactive elements of the interface. In addition, the system uses a number of HTML5 APIs [11], including web storage for persisting the user's system preferences.

The Samtla system libraries and data are encompassed by the *model* component. Samtla is written in Python and all system data is stored in SQL databases, except for the suffix tree, which is stored in *JSON* format and serialised to disk; see Section 4.2 for more detail on the suffix tree component.

The *model* is composed of a library of software tools that interact with the system data. The *search* component is responsible for answering user queries,

and uses a *Statistical Language Model* (SLM), which is a relatively recent framework for information retrieval [67]; see Section 4.1 for more detail on how we make use of SLM in Samtla. The SLM communicates with a suffix tree data structure, which is used to index the corpora that are being investigated. The suffix tree is loaded into memory at runtime for fast access. The suffix tree also provides support for the text mining tools, which are detailed in Section 5 and include a *related query* feature, which recommends queries to the user based on permutations of the original query resulting from morphological or orthographic variations present in the corpus (discussed in Subsection 5.1), a *related document* tool, presents users with a list of similar documents to the one they are viewing (discussed Subsection 5.2), and a *document comparison* tool facilitates the comparison of shared-sequences between documents (discussed in Subsection 5.3). Lastly, the *community* component is responsible for logging user data, such as query submissions and document views, usage statistics reflecting the user's navigation histories through the system, and for returning to the user recommended queries and documents based on their popularity in the user community; see Section 5.5 for more detail.

In the following sections we discuss in detail the methods and algorithms adopted to support the current set of tools divided into search, text mining, and community support.

4 Data Model

Statistical language modelling [67] is central to Samtla's data model. It provides the foundation for Samtla's search tool allowing users to locate documents through full and partial matches to queries. Samtla's *Statistical Language Model* (SLM), whose details are given in Subsection 4.1, is supported by a character-based *suffix tree* [36], described in detail in Subsection 4.2. A suffix trees is a very powerful data structure supporting fast retrieval of sequences of characters, known as *n-grams*, where n is the length of a sequence (for our purposes, measured in characters). Samtla is unique in that it is language agnostic and can thus support a variety of languages within a single data model; we will demonstrate this in Subsection 7.

4.1 Statistical Language Models

A SLM is a mathematical model representing the probabilistic distribution of words or sequences of characters found in the natural language represented by text corpora [58, 49, 67]. Samtla is designed as a language agnostic search tool and as such uses a character-based n-gram SLM, rather than the more conventional word-based model. Language modelling provide Samtla with a consistent methodology for retrieving and ranking search results according to the underly-

ing principles and structure of the language present in a corpus, which is often domain specific. Beyond that, SLMs provide a unifying model for Samtla's text mining tools described in Section 5.

Statistical language modelling combined with character-level (1-gram) suffix tree nodes enable the system to be applied to multilingual corpora with very little pre-processing of the documents, unlike word-based systems. For example, languages like Hebrew, Russian, and Italian attach affixes to a root word to identify syntactic relationships. This complicates word-based retrieval models since it is necessary to capture all instances of the same word in order to produce an accurate probabilistic model. Word-based models typically require a language-dependent stemming, part-of-speech tagging, or text segmentation algorithm, however, by adopting a character-based n-gram model these issues can be ignored to some extent and character-based models have been shown to outperform raw word-based models, especially when the language is morphologically complex [52]. Furthermore, a character-based model enables the system to be applied to different language corpora, but also corpora which contain documents written in several different languages. For example, some documents in the Aramaic collection contain texts written in Hebrew, Judeo-Arabic, Syriac, Mandaic, and Aramaic, the Vasari corpus contains English and Italian documents, whereas the British Library Microsoft Corpus covers a range of languages including English, French, Spanish, Hungarian, Romanian, and Russian.

Operationally, when the user submits a query, a list of documents is returned and ranked according to how relevant the document is to the query. The notion of relevance refers to the users expectation of which documents should be present at the top of the ranked list, in other words, which documents the user may be looking for [67]. In Samtla we take the view that the more probable a document in the SLM sense, the more relevant it is, thus avoiding the philosophical debate on the notion of "relevance" [55]. This equates to the system retrieving the most probable documents based on the SLM representing the distribution of the n-grams in the corpus being searched.

In Samtla the data comes from corpora, which consist of a collection of text documents grouped according to a specific topic, genre, demographic, or origin (e.g institution storing original versions of the digital texts). For instance, a Bible corpus can be composed of several Bibles from different periods or translations (Wycliffe Bible, Tyndale Bible, or Thomas Young's Translation). Each Bible therefore represents an individual corpus containing a collection of documents, which represent each chapter of the given Bible.

We generate SLMs from the corpora over the whole collection, which we call the *collection model*, and over each individual document, which we call the *document model*. A generic SLM is denoted by M, while the collection model is denoted by C and a document model is denoted by D. Each SLM is generated from the n-grams extracted from documents in the corpora, where n will vary from one

to some pre-determined maximum. The example below demonstrates how text is converted to n-grams of various sizes. Here the n-grams are generated for the sequence "beginning", which has a maximum n-gram size of nine, and can itself be reduced to lower-order n-grams by reducing the sequence a character at a time, as illustrated in the table below:

n−gram order	n−gram
9	beginning
8	eginning
7	ginning
6	inning
5	nning
4	ning
3	ing
2	ng
1	g

For the collection model, C, the (global) probabilities of the character-based n-grams are stored in a suffix tree, while for the document model, D, the (local) probabilities are stored in a conventional database to make them easily available for use by other system components. From an implementation perspective each document is represented by a unique *document ID* and the positions of the n-grams in documents are stored with the probabilities inferred from the language model. Thus, when a user submits a query, Samtla will compute the probability that the query was generated by the model, M, where M is either C or D. In other words, Samtla will compute the probability that a user who is interested in a given document in the collection would submit that query. The documents in the collection can then be ranked according to the computed probabilities for these documents and the top scoring documents are returned to the user.

We will now explain how the query model $P(q|M)$, denoted by $P_M(q)$ and read as "the probability that the query q was generated by the language model M", is computed; we will assume throughout that C represents the collection model. Using Bayes theorem [47]:

$$P(M|q) = \frac{P_M(q)P(M)}{P(q)} \tag{1}$$

The term $P(M|q)$ represents the conditional probability of the language model M given the query q. When M is a document D, this is the probability of the document D when the query is q, which will allow the system to rank the documents returned to the user. The right-hand side of (1) consists of the query model $P_M(q)$ multiplied by $P(M)$, the *prior* probability of the model M. When M is a document model D, the prior is its presupposed probability, which is often assumed to be uniform, i.e. the same for all documents and can thus be ignored for the purpose of ranking (see Section 9 where we discuss a non-uniform prior); Finally, the denominator $P(q)$, i.e. the probability of the query, is the

same for all documents and can thus also be ignored for the purpose of ranking. In summary, when the prior is uniform, we can rank documents according to $P_D(q)$, the query model for the document.

Let $q = c_1, c_2, ..., c_m$ be a sequence of m characters. Then, using the chain rule, the query model $P_M(q)$ is calculated as a product of conditional probabilities:

$$P_M(q) = P_M(c_1, c_2, ..., c_m) = \prod_{i=1}^{m} P_M(c_i|c_1, ..., c_{i-1}). \tag{2}$$

Each conditional probability, $P_M(c_i|c_1, ..., c_{i-1})$, on the right-hand side of (2), with $1 \leq i \leq m$, may be approximated by the maximum likelihood estimator (MLE) [51]:

$$MLE_M(c_i|c_{i-n+1}, ..., c_{i-1}) = \frac{\#(c_{i-n+1}, ..., c_i)}{\#(c_{i-n+1}, ..., c_{i-1})}, \tag{3}$$

where the # symbol before a sequence indicates its raw count in the model M. Moreover, for any sequence of characters we only make use of its n character history/context (or less than n for shorter sequences) as an approximation to the conditional probabilities in (2), in accordance with an n-order Markov rule [54]. (We also take $MLE_M(c_1|c_0)$ to be $MLE_M(c_1)$.)

A known problem with the maximum likelihood estimator (3) is when the raw count of a sequence is zero, resulting in estimating the query model probability, $P_M(q)$ in (1), also as being zero. This may be the result of the user entering a character or word incorrectly, for example spelling mistakes or typographical errors. Alternatively, the corpus may not be sufficiently large to encapsulate the full vocabulary of a given language, and thus due to this sparseness problem the query probability would be zero. To overcome this problem, *smoothing* [26, 69] adjusts the MLE probabilities to make them non-zero.

We smooth the MLE probability (3) via interpolation, using a weighted term, which defines the contribution to the overall probability for each order, k, where k varies from a zero order, 0-gram model, when $k = n + 1$, to an n-gram model, when $k = 1$. Each weight, represented by λ_k, defines the amount of interpolation, with lower-order models contributing less to the final probability.

Thus our approximation of the conditional probabilities on the right-hand side of (2) is given by the interpolation,

$$\hat{P}_M(c_i|c_1, \ldots, c_{i-1}) \approx \sum_{k=1}^{n+1} \lambda_k MLE_M(c_i|c_{i-n+k}, \ldots, c_{i-1}), \tag{4}$$

where we use $\hat{P}$ to make clear that we are approximating P, and the weighted

term for each k is given by

$$\lambda_k = \frac{n+2-k}{(n+1)(n+2)/2}, \tag{5}$$

where n is the order of n-gram, which is composed by interpolating the nth order model with lower order ones. When $k = n+1$, then $MLE_M(c_i|c_{i+1})$ is taken to be the 0-gram model, $\frac{1}{|V|}$, where V is the finite alphabet of the language (for English this is 26 representing the characters of the Roman alphabet).

If an n-gram of the query is not present for a particular document, then we need to back-off to a lower order n-gram. However, the MLE score for the lower order n-gram will be too high because lower order n-grams are often more frequent in a document than higher order n-grams. To reduce the influence of a missing n-gram on the final query score, we back-off to the lower order n-gram, and as before, extract the probability for the back-off n-gram given by the document model D and the collection model C (obtained by revisiting the suffix tree). Assuming we obtain a match for the back-off n-gram, we smooth the probability with a weighted normalisation term $\frac{n}{n+2}$ to provide a consistent normalisation for each order of n which defines the length of the n-gram obtained from the back-off procedure. If there is still no match, we store the result of $\frac{n}{n+2}$ and repeat the back-off process until we obtain a match for the lower order n-gram, or until we eventually arrive at the 0-gram model. The smoothed probability for the missing n-gram is then the sum of the probabilities obtained from the back-off and is used to approximate the conditional probability $\hat{P}_M(c_i|c_1, \ldots, c_{i-1})$ in Equation 4.

The next stage is to smooth the conditional probabilities. The historians we are working with tend to submit long and verbose queries representing "formulae", which can impact on retrieval performance [39], since they contain many uninformative query terms (such as prepositions: *of, in, to, by*, and determiners: *a, the*). To compensate for this we adopt the Jelinek-Mercer smoothing method, which involves a linear interpolation of the document model D with the collection model C using a coefficient represented by λ to control the influence of each model on the final query score. [68]. The final smoothed query score for a document is obtained by replacing $\hat{P}$ in (2) by P, as follows:

$$P_D(q) \approx \lambda \hat{P}_D(q) + (1-\lambda)\hat{P}_C(q), \tag{6}$$

where we chose $\lambda = 0.6$, i.e. 60% contribution from the document model to the smoothed query score. It is possible to further tune the value of λ by experimentation. As mentioned, long verbose queries require more smoothing than keyword or title queries due to the number of uninformative terms and so may require a higher setting for λ [68]. The further smoothing in (6) makes sense as even after the initial interpolation, the maximum likelihood document model probabilities may be low, while the maximum likelihood collection model probabilities will provide a better global estimate of the probability.

To summarise, the list of documents D containing the query are ranked according to $P_D(q)$ corresponding to the approximation in (6), which interpolates each document according to its language model and, in addition, interpolates the document and collection models. The assumption is that scoring documents in this manner presents to users the documents that are most likely to represent their information need. We further emphasise that the Jelinek-Mercer smoothing method has been adopted, since many of our users will submit long queries representing textual fragments (for example, a Bible verse), and this method of smoothing has been shown to be particularly effective for addressing long and verbose queries [68, 69]. In the next subsection we will describe the suffix tree, which provides the lower level implementation of Samtla's language model.

4.2 Suffix Tree

Samtla's search capability, based on SLMs as described in the previous subsection, is supported by a space optimised character-based suffix tree, with the aim of holding the complete data structure in memory for fast retrieval [36].

In order to reduce the suffix tree's memory consumption, we create a k-truncated suffix tree [61], which compresses the suffix tree by limiting its depth to k nodes at most, and store the data attached to tree nodes in an external key-value database. We have found that $k = 15$ works well for the languages we have experimented with, after plotting the length of words present in the corpus, which showed that the majority were no longer than 15 characters in length.

A further method to reduce the space requirement of the suffix tree is to compress dangling nodes, which are nodes that have only a single descendent (or child). These are gathered together during a depth-first-search and stored as a 'supernode', whose label is constructed from the concatenation of the collected node labels [36].

Given a text string, the resulting suffix tree represents a compressed "trie" data structure containing all the suffixes of the string as their keys and positions in the string as their values. A generalized suffix tree is a suffix tree constructed from a set of text strings instead of a single one, so its leaf nodes (represented by *$* in Figure 3) only need to store the ID and start position of the character string. In each node of the generalised suffix tree constructed from the corpus (document collection), we store the frequency of the corresponding string of characters, which is later used to perform the above-mentioned maximum likelihood estimations for the character-based n-gram SLM.

Once the generalised suffix tree is constructed, we calculate the conditional probabilities which form the basis for the collection model C (discussed in the previous section). The tree is traversed starting at the root node. For each node, the conditional probability is calculated by dividing the count of the cur-

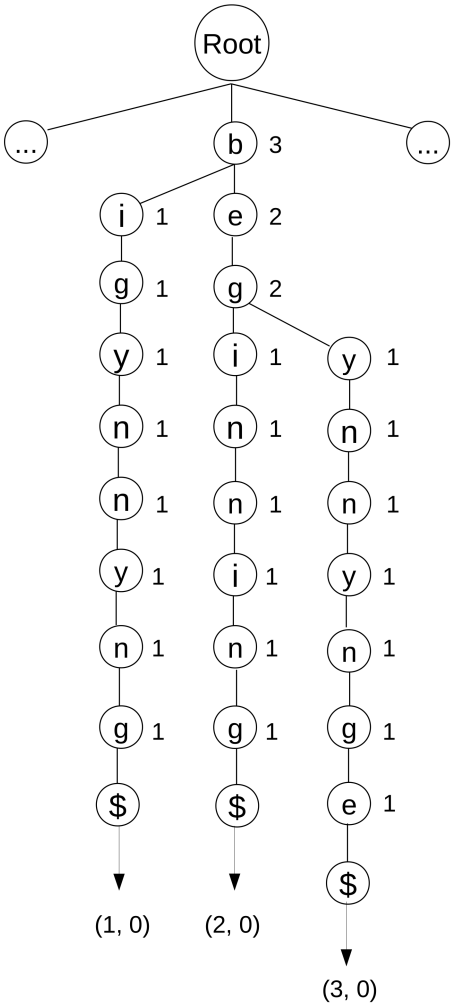

Figure 3: A truncated suffix tree over the strings "beginning', "bigynnyng', and "begynnynge"

rent node with the count of its parent node, as defined in Equation 3 above, and illustrated in Figure 5.

After calculating the initial probabilities the tree is traversed a second time to apply the interpolation between each order of n-gram (as detailed in Equation 4) to produce the final smoothed collection model C.

Searching a suffix tree is performed by starting at the root node and then descending the tree along a unique path by comparing characters of the query with

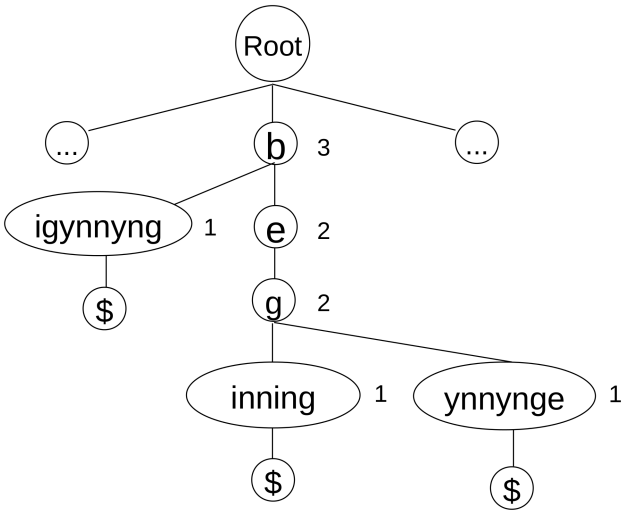

Figure 4: A compressed suffix tree where the super-nodes are rendered as ellipses

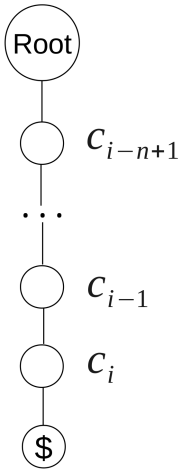

Figure 5: Calculating the MLE

the label stored at each node. When the characters of the query are exhausted or a mismatch occurs, the sub-tree rooted at the last-matched node is traversed with a breadth-first traversal, and all leaf nodes are collected resulting in an index of document IDs and start positions. Partial matches are also obtained during the traversal as we are always returning the last matched node, regardless of whether there are any further characters to match. (although see Section 5.1

for a discussion on how we can trigger different search strategies when there is a mismatch for the full query)

The language agnostic nature of the implementation has been tested with a number of corpora, which are both different in terms of structure, but also in terms of the language, dialect, script, and domain.

5 Text Mining Tools

Books, web pages, articles, and reports are all examples of unstructured text data where relevant information exists potentially anywhere within the document. Unstructured text data is often managed and retrieved via a search engine (see [46]). Search engines provide the means to retrieve information but not to analyse it, this is where text mining techniques are useful, as they provide the researcher with different views of the data that can enable them to discover and evaluate textual patterns [18].

The Samtla system is designed as a research environment packaged with a set of extensible text mining tools. The tools provide a means to analyse a corpus through the identification of patterns or "formulae" that are of potential interest to the researchers. In addition, the tools have been developed alongside each user group in order to identify the problem domain and provide solutions which can be implemented in accordance with the probabilistic approach adopted by the underlying system.

Each tool is built on one or more components of the data model (illustrated in Figure 2 and discussed in Section 4). For example, Samtla uses the collection model C and the suffix tree data structure to provide a related query tool, which generates and ranks queries similar to the users original search term (see Section 5.1). The related documents feature measures the similarity between document pairs using the language model D for each document, which is then ranked and presented to the user as a list of documents similar in content to the document they are viewing (see Section 5.2). This subset of tools fall under the recommendation component of the system, where statistical language models have been shown to perform well [43].

The tools are context-dependent as the system only presents the user with the results from a given tool if it makes sense in the given context. For example, if a user is viewing a document (document view), Samtla displays an informational side-bar containing related documents, selecting a document from the related documents list directs the user to an interactive document comparison tool (see Section 5.3) where shared-sequences can be compared between the two documents. Whilst in document view, the user has access to metadata for each document, which is tailored to each user group. In document view the user can also overlay additional data in the form of *named entities* [56], which are

labelled with additional metadata from external sources such as Wikipedia, online encyclopedias, and the Google maps API. The context-dependent design of the tools results in a minimal user interface, which dedicates more screen real estate to the data and tool output.

5.1 Related Queries Tool

The related queries tool extracts 10 text fragments from the corpus (specifically, from the suffix tree representing the collection model C as discussed in Section 4.1), that are most similar to the user's original query. These are then displayed as part of the ranked search results. For example, searching for "beginning" in the Bible edition of Samtla would return related queries "bigynnyng" and "begynnynge", which represent alternative spellings. The related queries are generated through a process similar to the Levenshtein edit distance algorithm [36] where alternative forms of the original query are created through processes of deletion, substitution, and insertion. These can be defined more formally as follows. Let Q represent the related query, where n is the length of the original query q, and $i \in 1, 2, \ldots, n$. Then the string edit processes are defined as:

Method	**Related Queries**
Deletion	$Q = q_1, q_2, ..., q_{i-1}, q_{i+1}, ..., q_n$
Substitution	$Q = q_1, q_2, ..., q_{i-1}, ?\ , q_{i+1}, ..., q_n$
Insertion	$Q = q_1, q_2, ..., q_{i-1}, ?\ , q_i, ..., q_n$

As an example, if the original query is "beginning", then the following related queries are generated.

Method	**Related Queries**
Deletion	eginning, bginning, ..., beginnig, beginnin
Substitution	?eginning, b?ginning, ..., beginni?g, beginnin?
Insertion	?beginning, b?eginning, ..., beginnin?g, beginning?

When the user submits a query, related queries are automatically extracted from the suffix tree component of the system. This is achieved by replacing each character of the original query one at a time with a wild-card character and then submitting them to the tree. As the wild-card character is not indexed by the suffix tree there will be a guaranteed mismatch at that point in the query. When a mismatch occurs, we execute the above functions, which traverse the suffix tree from the mismatched node, and attempt to match the remainder of the query. Deletion does not require a wild-card character since we simply remove a character from the string where the wild-card character would appear. Insertion and Substitution are achieved by replacing the wild-card character with the node label of each child rooted at the last matched node (its parent node), and attempt to match the remainder of the query after the

wild-card character. All successfully matched strings are returned as a ranked list of potential related queries according to their smoothed probability scores from the collection model C.

As illustrated above, the combination of the suffix tree data structure and the collection model component C of the language model provides a good basis for constructing a query recommendation system.

5.2 Related Documents Tool

Related documents are those documents that have common string sequences shared between them. The related documents tool finds up to twenty documents from the corpus that are most similar to the document currently viewed by the user (referred to as the target document). The retrieval of twenty documents is a user interface decision, since we compute similarity scores over all document pairs in the corpus. Each document, in the related documents menu, represents a link to the document comparison tool (discussed in Section 5.3).

A document can be considered a probability distribution over n-gram sequences [31], and the similarity between a pair of documents may be calculated through the Jensen-Shannon Divergence(JSD) over their corresponding document models M_d (as described in section 4.1) [48]. The JSD is the symmetric version of the well-known Kullback-Liebler Divergence (KLD) defined as:

$$D_{KL}(P||Q) = \sum_i P(i) \log_2 \frac{P(i)}{Q(i)} \ ,$$

where in our context P and Q represent two smoothed n-gram probability distributions provided by the corresponding document models, and i is a value drawn from the respective smoothed n-gram distribution based on a sliding window of size n. The smaller the sliding-window size n, the finer-grained the document similarity measure. In our experiments we have found that $n = 7$ provided a good balance based on our 15-gram language model. The JSD is calculated as

$$JSD(P||Q) = 1 - \sqrt{\frac{1}{2} D_{KL}(P||M) + \frac{1}{2} D_{KL}(Q||M)} \ ,$$

where M is the average of the two distributions $\frac{1}{2}(P + Q)$ [31]. The resulting JSD produces a score between 0 and 1, where a score of 1 means the documents are identical. The JSD scores are ordered in descending-order according to their similarity to the target document so that the most similar documents are ranked at the top of the related documents list.

5.3 Document Comparison Tool

A common task applied to language corpora is to find representative examples of language use exemplified by string patterns. Samtla provides a document com-

parison tool, where users can compare the document they are currently viewing with a document selected from a list of similar documents (discussed above in Section 5.2). This feature is considered to be essential by our user groups, since this form of comparison was performed manually and can be a complex task due to overlapping shared-sequences.

Although there exist some document comparison tools such as `diff` on UNIX, they perform *global sequence alignment*, which attempts to match the entire documents, while Samtla users are interested in *local sequence alignment* which identifies text regions of similarity within two documents that could be widely divergent overall. The underlying algorithm for identifying shared text patterns is a tailored variant of the Basic Local Alignment Search Tool (BLAST) algorithm [50], widely used in bioinformatics for comparing DNA sequences.

We first extract all trigram character strings shared by the given pair of documents as seeds, and then extend those seed strings one character at a time, first from the left and then from the right. During the iterative extension process, we score any pair of (approximately) matched strings s_1 and s_2 by their Levenshtein edit distance [36], which measures the number of changes required to convert one string in to another using deletion, insertion, and substitution. The metric is defined as:

$$ed(s1, s2) \leq \lfloor m\delta \rfloor \tag{7}$$

The extension stops when the edit distance reaches the floor of a certain threshold $m\delta$, where m is the length of the shorter matched string between s_1 and s_2, multiplied by a tunable tolerance parameter δ. The default setting is $\delta = 0.2$ which is equivalent to a 20% difference between the two sequences, before moving on to the next seed. Characters representing punctuation are ignored during the extension process.

As an example, a text pattern found by Samtla in the Bible corpora, starting from the trigram seed string *ham*, is shown as follows.

King James Bible: Noah; Shem, **Ham**, and Japheth
Douay-Rheims Bible: Noe: Sem, C**ham**, and Japheth

The above example has a total edit distance of 4. The strings *Noah* and *Noe* have an edit distance of 2, since there is one substitution ($a \rightarrow e$) and one deletion (final character h) required to convert *Noah* to *Noe*. The strings *Shem* and *Sem* require one deletion of character H which is equal to an edit distance of 1, and likewise, *Ham* and *Cham* is converted with the insertion of character C at the beginning of the string.

It can be seen that our method captures text patterns, which differ in terms of spelling errors or orthographic variations. Despite such superficial differences, a

researcher of Bible scripture would probably consider those two text fragments as identical passages (Chapter 10, Genesis). For a discussion of the document comparison interface see Section 6.4.

5.4 Named Entity Tool

Named Entity Recognition (NER) [22] describes the process of extracting words (or sequences of characters, in our case), that represent names of people, companies, and locations. Samtla uses gazetteers to extract named entities from the raw documents. Gazetteers have been used for some time to improve the performance of named entity systems, other more sophisticated methods exist, for instance, semi-supervised learning techniques such as bootstrapping [18], however gazetteers are becoming popular once again due to the wealth of structured data on named entities provided by platforms such as Wikipedia and DBpedia [41]. A further motivation for adopting the gazetteer approach, is that the current versions of Samtla support a number of historic text collections, such as the Bible and Vasari's 'the lives of the most excellent artists and architects' 7, these collections represent closed corpora, which means that there are rarely going to be new documents added to the collection. Consequently, gazetteers are sufficient for these types of domain specific and static corpora as there are a wealth of lists already compiled by researchers that can be used to form the basis for gazetteers. Furthermore, we have found that Wikipedia can be leveraged since there are a large number of general lists of people, locations, and other miscellanea [15], but also lists for specific collections [14, 13].

Named entities are located in the documents by submitting each entry in the gazetteer to the suffix tree as queries. Each full match is stored in a database organised by entity type together with the document identifier and an index of start positions in the text. The data is parsed to the browsing tool (see Section 6.1), which provides further entry points to the documents, with the named entities themselves being rendered as an additional layer over the document in document view (further discussed in Section 6).

The gazatteers could also be used to form the basis of training data for a statistical learning approach [33] to enable Samtla to identify and mark up documents semi-automatically, which is an approach that we will be investigating as part of future work.

5.5 Recommendation Tools

Personalised recommendation systems are familiar to many users of the internet. For instance, online shoppers often encounter the 'what other customers bought' page, which presents a series of recommended items that other buyers purchased based on items in their 'basket' or 'shopping cart'. Other examples of recommendation systems can be found in socially affected personalisation where a user is part of a select group who share content and opinions with other users

they trust, as well as collaborative search, which enables users to discover new search terms based on the search behaviours of other users [65]. Samtla leverages user activity to generate recommended queries and documents. Through analysing the log data, Samtla can inform users of the top-10 most popular queries and documents in the research community, so as to support users' collaborative search. Thus a user of Samtla can be directed to the "interesting" aspects of the corpus being studied, which may not have occurred to them previously.

Log files are used to store usage statistics, user interaction through the system using referrer URLs, and system error reports. This data can be leveraged in interesting ways, one of which is to return the users search and page view history. Users may also wish to discover what is popular in a corpus, as a way to find new documents of potential interest. The current version of Samtla supports a community feature which suggests search terms and document views based on their popularity, this requires storing data such as unique userIDs, timestamps, queries and document IDs. The user data is then used to produce top-ten ranked lists of queries and document views per user and the community as a whole.

The popular queries and documents are ranked and selected using an algorithm similar to the Adaptive Replacement Cache (ARC) [53], where the frequency of each query or document is combined with its recency (measured by the number of days that have passed since the last submission of the query or document view), and used for ranking. This ensures that the recommended popular queries and documents are biased towards fresh ones and updated along with time. Formally the popularity of a query or document is defined as

$$popularity = T^{\beta} R^{1-\beta} \tag{8}$$

where $T = \frac{1}{S}$ with S representing the count in days since the last submission, where $today = 1$, with weighted term $\beta = 0.6$, and parameter R which represents the raw count of submissions for the query or document. The combination of the two terms T and R prevent submissions with high counts, but longer time between submissions, from dominating the top entries of the recommended queries or documents.

The resulting ranked results are made accessible via the respective side-bars in the user interface, which are populated when the user navigates to a document through browsing or searching (see Section 5.4).

6 User Interface

Interaction with the system is through a browser-based interface, the user can perform three main tasks in the current version of the system: (i) corpus brows-

ing (see Section 6.1), (ii) search (Section 6.2), and (iii) document comparison (Section 6.4).

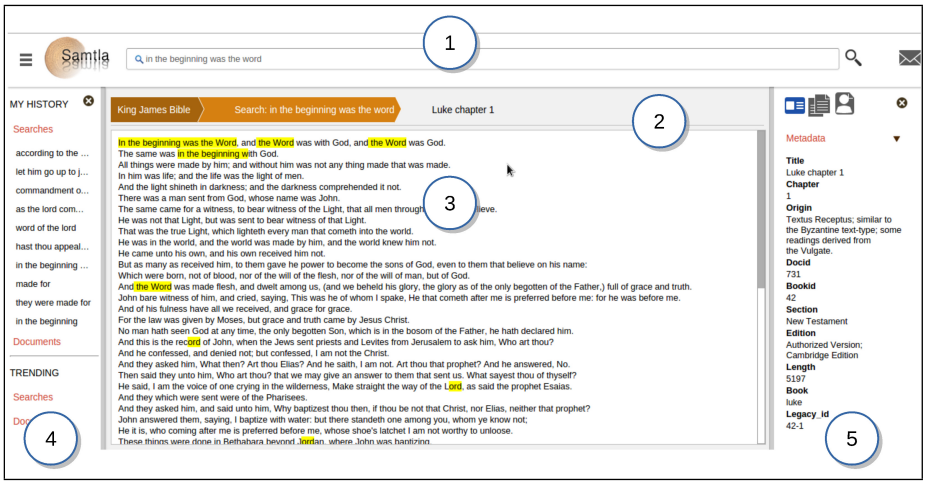

Figure 6: Samtla User Interface showing 1.) the search bar, 2.) breadcrumb-based navigation 3.) main window, which displays search results, and the document text with additional information such as highlighted query terms (shown here), and output of the text mining tools, 4.) the query and document view history for the user and community (trending searches and documents), and 5.) a side-panel for displaying the metadata, related documents (for accessing the document comparison tool), and activating additional data layers over the document e.g. named entities.

The tool set is designed to be modular and extensible in order to enable further tools to be developed with our users without affecting previously established system components (see Sections 3, 7, and 9).

6.1 Browse Documents

Samtla adopts a clustered (or faceted) navigation model, where each cluster describes a category represented by a collection of documents sharing a common property [24] [28]. Clustering documents according to a particular feature [63], can provide users with an indication of the type and availability of data in a system [38], and is a useful approach for encouraging users to explore and discover information within a collection [37]. By adopting a clustered navigation model, future components can be integrated in a modular fashion, without introducing visual clutter through traditional UI elements like tabs and drop-down menus.

The browsing architecture is divided in to two separate presentation layers. The default is a list view, which mimics a traditional file directory where each row entry represents either a folder or an individual document. Columns con-

tain the cluster label or document name, and further information extracted from the document metadata (see Figure 7). The alternative view uses a squarified

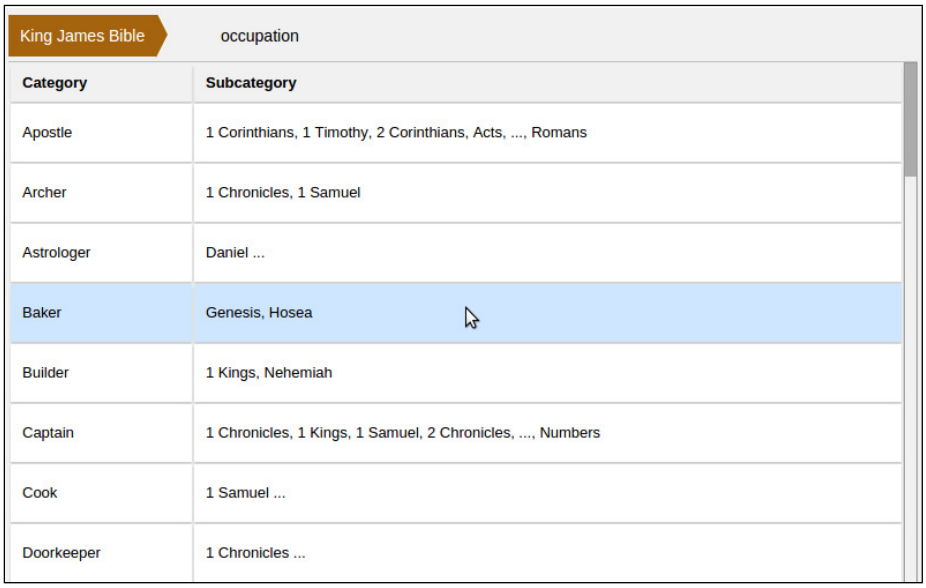

Figure 7: Browsing the Bible corpus using the list view

treemap [25], and can be considered as representing a topic model (see Figure 8), where each topic provides the user with a different clustered view of the corpus generated from the metadata or the named entity tool (see Section 5.4). The user can switch between the list and treemap views via a button in the interface, as users may prefer one form of presentation over another.

The advantage of the treemap representation is that it is very flexible and can be enriched with textual or visual information by populating the cells with metadata or by altering the dimensions or colour of the cells to indicate membership or extent. For instance, the size of each cell can be adjusted to reflect that a document is longer or that a cluster contains more members than others or a preview image of the document could be displayed to aid navigation.

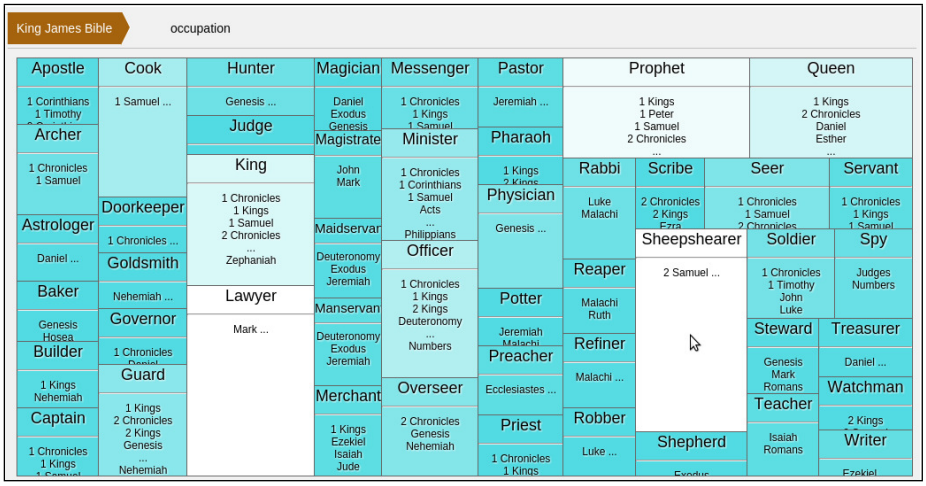

Figure 8: Browsing the Bible corpus through the treemap view generated from the document metadata

6.2 Search Documents

Query results in Samtla are divided in to two types: exact and partial matches. Partial matches do not encompass the full query, in other words, not all characters of the original query were matched. However, partial matches are returned, since they could still be of interest to the researcher, but will appear lower down in the search results below full query matches. For example, if we consider the query "*pharaoh was wroth against his two officers against the chief of the butlers*", from the Bible, Samtla will return (aside from exact matches), examples of other roles exemplified by the inexact matches "*the chief of*", such as; "*chief of the cup-bearers*", "*chief of the bakers*", "*chief of the tower/round-house*", and "*chief of the eunuchs*", which tells us something about roles within the King's court at that time (see Figure 9). Alternatively, the results returned by a partial match may help the researcher to reformulate their query given them a basis to start from - in contrast to boolean forms of search, which return no results if the exact query was not found.

When a search is performed, we obtain an index from the suffix tree (see Section 4.2) containing the document id and the start and end positions of the matched n-grams for each document. This is passed to the language model for ranking the documents according to the query, and a snippet generation tool for producing short snippets of the match query for the search result view. We order the ranked list by sorting the results first according to the length of the matched query, and then by the probability of the query given the query score for the document, retrieved from the SLM (see Section 4). The top ranks of the results reflect the highest scoring or most relevant documents.

Each document in the search results is then rendered with its title and the generated snippet window showing the preview of the query in the document, which contains the top-3 snippets that best describe the query. Snippet windows were selected as the most appropriate method for summarising the document, as they are familiar to users [38].

The snippets enable the user to evaluate the relevance of each document in the ranked list before deciding which document best meets their information need. Snippet length is tunable and we define a parameter w, which limits the maximum length of each snippet, for our purposes this is set to $w = 100$ characters, however in future versions this could be provided as a user setting.

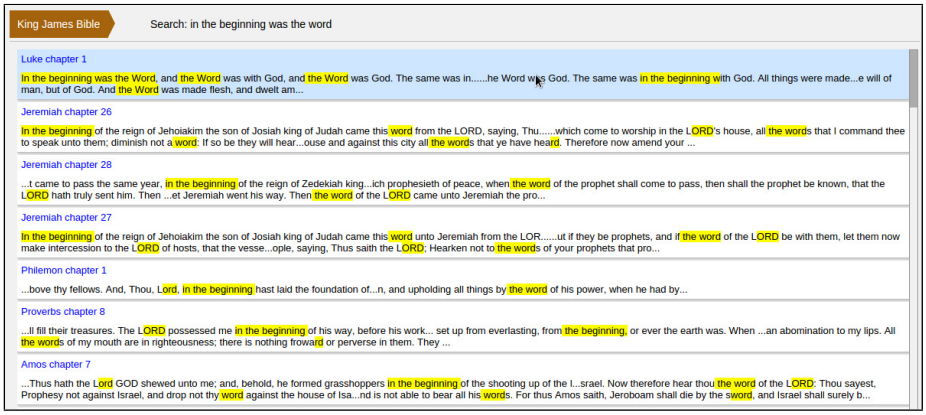
King James Bible
Search: in the beginning was the word
Luke chapter 1
In the beginning was the Word, and the Word was with God, and the Word was God. The same was in.....he Word was God. The same was in the beginning with God. All things were made...e will of man, but of God. And the Word was made flesh, and dwelt am...
Jeremiah chapter 26
In the beginning of the reign of Jehoiakim the son of Josiah king of Judah came this word from the LORD, saying. Thu.....which come to worship in the LORD's house, all the words that I command thee to speak unto them; diminish not a word: If so be they will hear...ouse and against this city all the words that ye have heard. Therefore now amend your ...
Jeremiah chapter 28
...t came to pass the same year, in the beginning of the reign of Zedekiah king...ich prophesieth of peace, when the word of the prophet shall come to pass, then shall the prophet be known, that the LORD hath truly sent him. Then ...et Jeremiah went his way. Then the word of the LORD came unto Jeremiah the pro...
Jeremiah chapter 27
In the beginning of the reign of Jehoiakim the son of Josiah king of Judah came this word unto Jeremiah from the LOR.....ut if they be prophets, and if the word of the LORD be with them, let them now make intercession to the LORD of hosts, that the vesse...ople, saying, Thus saith the LORD: Hearken not to the words of your prophets that pro...
Philemon chapter 1
...bove thy fellows. And, Thou, Lord, in the beginning hast laid the foundation of...n, and upholding all things by the word of his power, when he had by...
Proverbs chapter 8
...ll fill their treasures. The LORD possessed me in the beginning of his way, before his work... set up from everlasting, from the beginning, or ever the earth was. When ...an abomination to my lips. All the words of my mouth are in righteousness; there is nothing froward or perverse in them. They ...
Amos chapter 7
...Thus hath the Lord GOD shewed unto me; and, behold, he formed grasshoppers in the beginning of the shooting up of the l...srael. Now therefore hear thou the word of the LORD: Thou sayest, Prophesy not against Israel, and drop not thy word against the house of Isa...nd is not able to bear all his words. For thus Amos saith, Jeroboam shall die by the sword, and Israel shall surely b...

Figure 9: The search result interface - showing ranked documents with snippet windows

The snippet scoring algorithm extracts all potential snippets and ranks them by interpolating the length of the set of ngrams in the query found in the snippet with the total count of all terms appearing in the snippet, where more weight is assigned to snippets containing all of the query terms. This ensures that the snippet window is ranked in such a way that the top snippets will contain all of the terms of the users' query before presenting snippets with only partial matches of the query. Let δ be the cardinality of the set of n-grams that are present in the snippet, which are present in the query and μ be the count of all query n-grams (including repetition) found in the snippet, with a weighted term $\alpha = 0.9$, which as mentioned, biases the snippet towards those that contain all parts of the query, then each snippet can be scored through,

$$snippet\ score = \delta^{\alpha}\mu^{(1-\alpha)}, \tag{9}$$

The snippets are then sorted in descending order by the score returned in Equation 9, and the top-3 selected as a preview for the document. When the user

selects a snippet the system opens the document and scrolls to the location of the selected snippet. Other components of the search interface include the related queries, which are displayed above the search results as links ordered by their probability given the collection model C (see Section 4), from left-to-right.

6.3 View Documents

When a user arrives at the document level through browsing or searching, they are presented with a main window displaying the document text, or where available the image of the scanned document, see Figure 6. If the user has navigated to the document through the browsing tool, then any metadata related to the document, including the named entity (see Figure 10, is highlighted, for documents located using the search tool, we highlight all instances of the query. In document view, the user has access to the metadata, document comparison, and named entity tools.

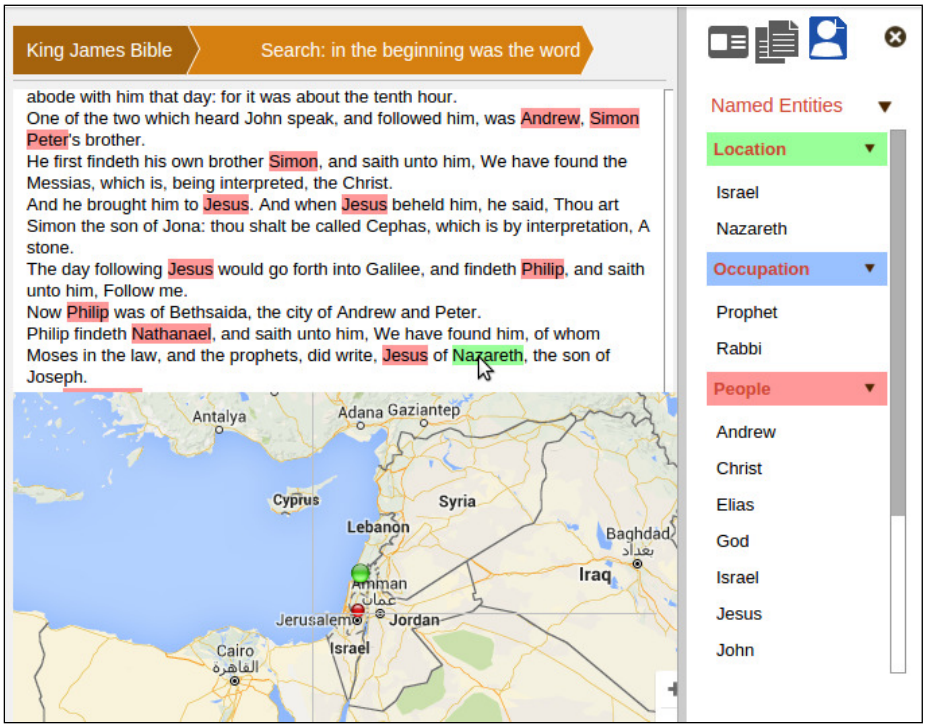
King James Bible
Search: in the beginning was the word
abode with him that day: for it was about the tenth hour.
One of the two which heard John speak, and followed him, was Andrew, Simon Peter's brother.
He first findeth his own brother Simon, and saith unto him, We have found the Messias, which is, being interpreted, the Christ.
And he brought him to Jesus. And when Jesus beheld him, he said, Thou art Simon the son of Jona: thou shalt be called Cephas, which is by interpretation, A stone.
The day following Jesus would go forth into Galilee, and findeth Philip, and saith unto him, Follow me.
Now Philip was of Bethsaida, the city of Andrew and Peter.
Philip findeth Nathanael, and saith unto him, We have found him, of whom Moses in the law, and the prophets, did write, Jesus of Nazareth, the son of Joseph.
Antalya
Adana
Gaziantep
Cyprus
Syria
Lebanon
Baghdad
بغداد
Iraq
Amman
عمان
Jerusalem
Jordan
Israel
Cairo
القاهرة
Named Entities
Location
Israel
Nazareth
Occupation
Prophet
Rabbi
People
Andrew
Christ
Elias
God
Israel
Jesus
John

Figure 10: The Bible version of Samtla showing the document view with the named entity layer.

6.4 Compare Documents

Related documents (see Section 5.2), provide access to the document comparison tool. The document comparison tool is composed of two document windows, one for the target document (the document the user is currently viewing), and another for the document selected from the list of related documents. Each time a user selects a new related document, both documents are updated with new sequence data and the longest shared-sequence is highlighted in each as a starting point for the user. The tool is equipped with a control to choose

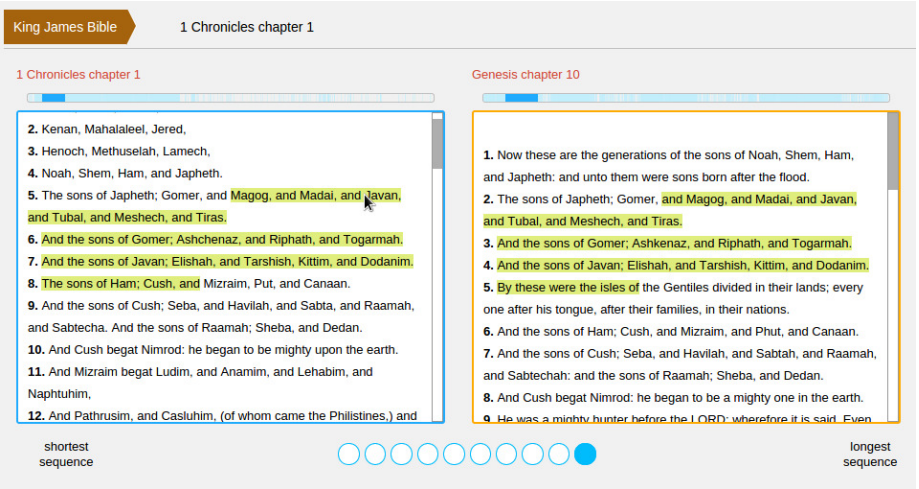
King James Bible
1 Chronicles chapter 1
1 Chronicles chapter 1
2. Kenan, Mahalaleel, Jered,
3. Henoch, Methuselah, Lamech,
4. Noah, Shem, Ham, and Japheth.
5. The sons of Japheth; Gomer, and Magog, and Madai, and Javan, and Tubal, and Meshech, and Tiras.
6. And the sons of Gomer; Ashchenaz, and Riphath, and Togarmah.
7. And the sons of Javan; Elishah, and Tarshish, Kittim, and Dodanim.
8. The sons of Ham; Cush, and Mizraim, Put, and Canaan.
9. And the sons of Cush; Seba, and Havilah, and Sabta, and Raamah, and Sabtecha. And the sons of Raamah; Sheba, and Dedan.
10. And Cush begat Nimrod: he began to be mighty upon the earth.
11. And Mizraim begat Ludim, and Anamim, and Lehabim, and Naphtuhim,
12. And Pathrusim, and Casluhim, (of whom came the Philistines,) and
Genesis chapter 10
1. Now these are the generations of the sons of Noah, Shem, Ham, and Japheth: and unto them were sons born after the flood.
2. The sons of Japheth; Gomer, and Magog, and Madai, and Javan, and Tubal, and Meshech, and Tiras.
3. And the sons of Gomer; Ashkenaz, and Riphath, and Togarmah.
4. And the sons of Javan; Elishah, and Tarshish, Kittim, and Dodanim.
5. By these were the isles of the Gentiles divided in their lands; every one after his tongue, after their families, in their nations.
6. And the sons of Ham; Cush, and Mizraim, and Phut, and Canaan.
7. And the sons of Cush; Seba, and Havilah, and Sabtah, and Raamah, and Sabtechah: and the sons of Raamah; Sheba, and Dedan.
8. And Cush begat Nimrod: he began to be a mighty one in the earth.
9. He was a mighty hunter before the LORD: wherefore it is said, Even
shortest sequence
longest sequence

Figure 11: An example of Samtla document comparison. The document comparison interface shows a pairwise comparison of the target document (left) and a document selected from the list of related documents (right). Sequences highlighted in yellow reflect the currently selected sequence, and blue represents all sequences shared between the two documents.

the length of the shared text pattern to view, with the minimum being 3-gram and the default setting displaying the longest sequence found between the two documents. This enables users to investigate large shared-sequences spanning several lines to smaller sequences representing a word or grammatical affix (typically 3-gram in length). Appearing above each document is a small horizontal map summarising all shared-sequences in the document, which provides the user with an overview of how the sequences are distributed throughout the two documents. For example, the shared-sequences may all appear in the introduction or abstract of the text. Clicking on a shared-sequence in a document highlights all instances of that sequence across both documents (see Figure 11). Sequence comparison is difficult to perform manually, especially over several documents and particularly when some of the sequences may be approximate, or overlap one another. The design of the document comparison tool is based on feedback from our users, and the orientation of the document windows attempts to emu-

late the manual process of document comparison, where a user may layout two documents, or pages side-by-side.

7 Case Studies

There are currently five versions of Samtla, with two versions serving two separate groups of digital humanities researchers. The first user group is represented by a team of historians led by the University of Southampton[7] who are analysing a corpus of 650 Aramaic Magic Bowls and Amulets from Late Antiquity (6th to 8th CE) written in a number of related dialects including Aramaic, Mandaic, and Syriac. The texts are written in ink on clay bowls, and cover a wide subject matter. The research involves searching and comparing textual fragments, which are formulaic in nature and provide an insight into the development of liturgical forms which differ due to transmission over centuries, and orthographic variation as a result of differences in authorship or dialect. There are also transcription errors resulting from damage to the original artefact, or illegible characters. Existing tools were not sufficient for identifying approximate text fragments meaning the analysis was largely a manual process of comparison and documentation.

The second user group is the Vasari Research Centre. The documents represent chapters from the book *Lives of the Most Excellent Painters, Sculptors, and Architects* by Giorgio Vasari (1511 - 1574). Giorgio Vasari is considered to be the founding father of the Art History discipline [17]. The Vasari Samtla contains documents in the original Italian and a corresponding English translation, and is used for research and as a teaching aid for students in class. Users can view either version by searching, browsing, or selecting the alternative version from the metadata in document view, which is then displayed side-by-side for comparison (see Section 6 for more detail on the User Interface). The Vasari corpus also contains a large number of images of paintings and architecture, which are displayed with the document metadata.

A third version of Samtla is applied to the Microsoft corpus of 68,000 scanned documents, which was bequeathed to the British Library. The collection represents books digitised from the worlds libraries and contains a range of languages, literary genres, over a couple of centuries. Moreover, as a proof-of-concept, a special edition of Samtla has been applied to the King James Bible, in English. This version is used for demonstration purposes and evaluation as many people are familiar with the content of the Bible.

The most recent Samtla was constructed for a pilot study between the British Library and the Financial Times (FT). The documents are represented by a corpus of Newspapers that have been digitised using OCR technology. The OCR data was provided along with the scanned pages of the Newspaper, which cover the year 1888, 1939, 1966, and 1991. This particular archive, required new tools

that could leverage the image data in order to compensate for poor quality OCR that reflected the current state of the art at the time of digitisation. Much of the text for the earlier articles (e.g. 1888, 1939, and to some degree 1966) are not reliably searchable due to poor recognition rates, and consequently the focus was on developing a metadata search component to complement the existing search tool, allowing users to search both the metadata and the full document. In addition, this Samtla presents users with the original image (see Figure 12), which utilises the document metadata to render boundaries around the articles and to make them selectable so that users can navigate the articles contained in a single newspaper.

Figure 12: The FT version, with the document view showing the named entity layer rendered on top of the original image.

The user is able to navigate between the raw OCR text and the scanned image, which required that some existing tools required some adaption to make use of the image data. For instance, the named entity tool (see Section 5.4) was originally developed for text data, but the FT version renders the named entities in both the raw text and the original scanned image, providing the ability to select and filter named entities in both views.

Each version of Samtla differs only in terms of the respective document collection, while the underlying system remains unchanged due to its language-independent and data-driven design. Upgrades and new features are rolled out across all versions, meaning that all user groups benefit from tools developed in collaboration with each user group.

8 Evaluation

8.1 Overview

In this section we describe the evaluation process for measuring the performance of the Statistical Language Model underlying the Samtla search engine (see Section 4). The evaluation assesses the ranking quality of the Samtla search engine in terms of whether the system (see Section 4) consistently puts the most relevant documents at the top positions of the search results list.

8.2 Crowdsourcing

Crowdsourcing is a web-based business model [23] that enables companies and individuals to employ the skills of people from a distributed community in order to perform some task in return for a small reward. These tasks are often large in scale or complex, and therefore time consuming as a result.

Crowdsourcing in Information Retrieval has generally involved outsourcing manual tasks such as data-annotation, labeled-data collection for training models, and system evaluation. This process was often completed in-house with a limited workforce, which depending on the size of the task, could be a slow process involving several days of work. Due to the size of the crowd, who are globally dispersed, tasks can be completed much faster at any hour of the day. There is also the potential for reducing bais in aggregated results, compared to in-house evaluations, due to the diversity and representativeness of the workers in terms of demographic [44].

There are a number of crowdsourcing platforms available for running surveys and evaluations. Amazon Mechanical Turk[1] (*MTurk*) is one of the better known ones [42, 19, 66], however, it is only available to researchers resident in the United States of America. As a result we selected Prolific Academic[2] [6], a crowdsourcing platform for academics and part of the Software Incubator at the University of Oxford. Prolific Academic currently have a participant pool of over 22,000 participants (as of 27/12/2015). The platform directs users to a website hosting a static survey or application. When the user completes the survey, they are presented with a URL, which activates a payment for their completed submission.

The majority of crowdsourcing platforms we investigated provided support only for static surveys, where the evaluation is represented by a series of static web pages constructed using a template web form editing tool. Platforms that allow the researcher to link to a URL hosting a web application, provide more flexibility by enabling the evaluation software to perform some action or logic based

[1] `https://www.mturk.com/`
[2] `https://www.prolific.ac/`

on user input, including monitoring the quality of the results, or distributing groups of tasks across different groups of users. Like MTurk, Prolific Academic directs users to an external website hosted by the researcher, which makes it possible to implement a survey tool that can serve dynamic content to the users, and also record information in the form of log files during the evaluation.

8.3 Methodology

The evaluation consisted of 50 queries represented by a ranked list of the top-10 documents returned by the Samtla system. The users were asked to assign graded relevance scores according to the four relevance grades "not relevant", "somewhat relevant", "quite relevant", or "highly relevant" to each document in the ranked list based on a given query displayed at the top of each search result.

8.3.1 Data Preparation

To evaluate the ranking performance of the system we used the King James Bible version of Samtla, since many people are familiar with the content of the Bible to some degree. We prepared a set of 50 queries of variable length, ranging from single word queries (i.e. "Moses", and "Jesus Christ"), to longer verbose queries representing set phrases (i.e. "the Lord hath spoken", and "blessed be the Lord"). We also constructed two test queries in order to have some control over the quality of users.

Each query was submitted to the Samtla search engine and the documents for the top-10 results were selected. The queries are processed to create two permutations of the ordering of the documents. The first permutation is a ranked list where the documents are sorted by their Statistical Language Model (*SLM*) score, which we will label as the *Samtla* order queries. The second permutation is generated by shuffling the position of the documents, which we refer to as the *random* order queries. Each user completed 10 queries in the *Samtla* order, and the remaining 40 queries in *random* order.

We measure system performance using the *random* order queries exclusively. The documents are sorted by their *SLM* score to recreate the *SLM* ranking, which we compare to the user-generated ranking, which we call the *consensus* ranking. The *consensus* ranking is created by aggregating all users' relevance grades for each document where "Not relevant"$= 1, ...,$ "Very relevant"$= 4$. We test for a presentation bias [20] by comparing this *consensus* ranking to the *display* order of the documents using both the full set of 50 queries and the 40 *random* order queries. If users are influenced by the presentation order of the documents, we will find documents at the top of the *display* ranking being assigned higher relevance grades simply due to their position in the ranked list, which may actually be in lower positions according to the *SLM* ranking. A display bias will be apparent if there is a notable difference in the average scores

across the two query sets. When discussing the performance measures we let r_1 and r_2 denote the *system* ranking and *consensus* ranking respectively.

8.3.2 User Interface

The evaluation interface represents a cut-down version of the Samtla system, which isolates the search result window. At the top of the page we display the query submitted to the system, which was used to generate the list of top 10 search results. Each document in the ranked list is displayed with the title and a short snippet showing the top three fragments containing the highlighted query in the document. Alongside each document is a drop-down box where the user selects an appropriate relevance grade.

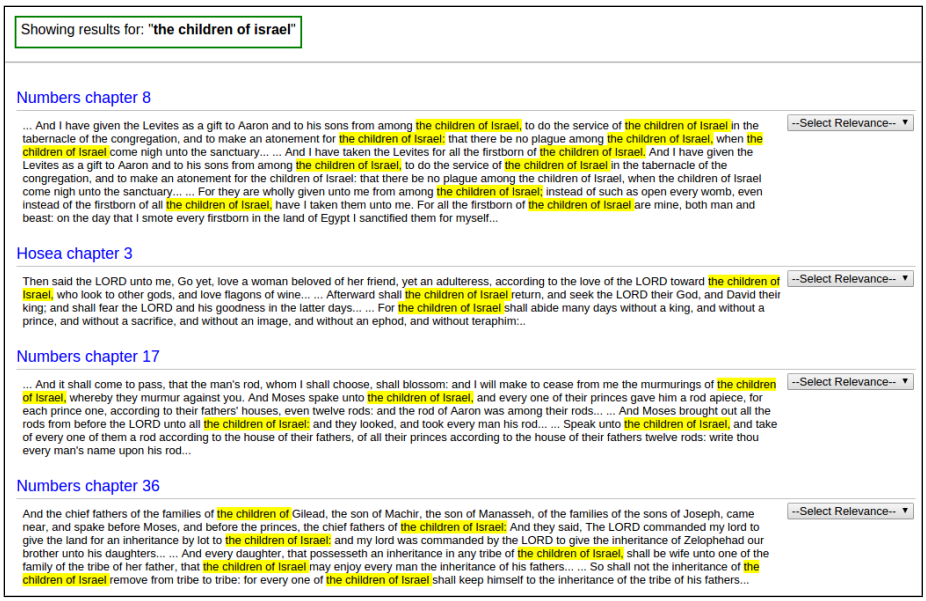
Showing results for: "the children of israel"
Numbers chapter 8
... And I have given the Levites as a gift to Aaron and to his sons from among the children of Israel, to do the service of the children of Israel in the tabernacle of the congregation, and to make an atonement for the children of Israel: that there be no plague among the children of Israel, when the children of Israel come nigh unto the sanctuary... ... And I have taken the Levites for all the firstborn of the children of Israel. And I have given the Levites as a gift to Aaron and to his sons from among the children of Israel, to do the service of the children of Israel in the tabernacle of the congregation, and to make an atonement for the children of Israel: that there be no plague among the children of Israel, when the children of Israel come nigh unto the sanctuary... ... For they are wholly given unto me from among the children of Israel; instead of such as open every womb, even instead of the firstborn of all the children of Israel, have I taken them unto me. For all the firstborn of the children of Israel are mine, both man and beast: on the day that I smote every firstborn in the land of Egypt I sanctified them for myself...
--Select Relevance--
Hosea chapter 3
Then said the LORD unto me, Go yet, love a woman beloved of her friend, yet an adulteress, according to the love of the LORD toward the children of Israel, who look to other gods, and love flagons of wine... ... Afterward shall the children of Israel return, and seek the LORD their God, and David their king; and shall fear the LORD and his goodness in the latter days... ... For the children of Israel shall abide many days without a king, and without a prince, and without a sacrifice, and without an image, and without an ephod, and without teraphim:..
--Select Relevance--
Numbers chapter 17
... And it shall come to pass, that the man's rod, whom I shall choose, shall blossom: and I will make to cease from me the murmurings of the children of Israel, whereby they murmur against you. And Moses spake unto the children of Israel, and every one of their princes gave him a rod apiece, for each prince one, according to their fathers' houses, even twelve rods: and the rod of Aaron was among their rods... ... And Moses brought out all the rods from before the LORD unto all the children of Israel: and they looked, and took every man his rod... ... Speak unto the children of Israel, and take of every one of them a rod according to the house of their fathers, of all their princes according to the house of their fathers twelve rods: write thou every man's name upon his rod...
--Select Relevance--
Numbers chapter 36
And the chief fathers of the families of the children of Gilead, the son of Machir, the son of Manasseh, of the families of the sons of Joseph, came near, and spake before Moses, and before the princes, the chief fathers of the children of Israel: And they said, The LORD commanded my lord to give the land for an inheritance by lot to the children of Israel: and my lord was commanded by the LORD to give the inheritance of Zelophehad our brother unto his daughters... ... And every daughter, that possesseth an inheritance in any tribe of the children of Israel, shall be wife unto one of the family of the tribe of her father, that the children of Israel may enjoy every man the inheritance of his fathers... ... So shall not the inheritance of the children of Israel remove from tribe to tribe: for every one of the children of Israel shall keep himself to the inheritance of the tribe of his fathers...
--Select Relevance--

Figure 13: The evaluation page showing a single test.

8.3.3 Selecting participants

Prolific Academic provides a number of filters to enable researchers to exclude certain users based on specific attributes stored as part of the users profile. The main filtering criteria applied to our study was to ensure that users were fluent English speakers. When we speak of a participant it should be clear from the start that they represent a member of the public and are not necessarily concerned with the motivation behind the specific study, or that they have a background in the type of data you are presenting to them as a researcher. Therefore it is important to prepare for this fact and attempt to filter the crowd of individuals for those who will be competent in completing the required task.

As mentioned above, the evaluation contained two test queries at the beginning of the survey. Not all users will have read or understood the instructions [34], and others may simply assign relevance grades at random in order to complete the survey and receive payment as quickly as possible, known as "gaming" the system [42]. It is important to plan and mitigate against these types of user behaviour, especially when it comes to crowdsourcing, since it is generally not feasible to monitor the performance of users in realtime during the evaluation (although see [66]).

The first test query contained the top-5 ranked documents for the single word query "Satan" displayed at the top of the page. The remaining 5 entries of the search results contained the snippets from a completely different, much longer query, "chief priests and scribes". To pass the test, the user has to assign "Not Relevant" to these last five documents since they do not match the query "Satan".

The second test query "Jesus Christ" was composed of the top-10 documents ranked in reverse order of relevance. In order to continue on to the evaluation, the user must assign higher relevance grades to documents as the rank position increases. The results presented in the next section demonstrate that test queries are an important design consideration.

8.4 Evaluation Measures

We adopt two sets of measures for calculating the system performance. The first set of measures assesses the correlation between the system and the user generated ranking over each query in the *SLM* rank and *display* rank order. If the system ranking is highly correlated with the user ranking then we can conclude that the ranking performance of the system closely matches that of a human assessor. The second measure evaluates the ranking quality of the system using the Normalised Discounted Cumulative Gain measure (NDCG), which is commonly adopted for graded-relevance based evaluations [40]. We perform the measures over both *SLM* and *display* permutations. In the following subsections, we describe each of the measures in more detail, before presenting a summary of the final results.

8.4.1 Correlation Measures

We measure the degree of correlation between the system and the users with *Spearman's footrule* [30] and the *M*-measure variant [21]. These non-parametric measures describe the degree of correlation between two ranked lists, and provide similar results to other correlation measures including Spearman's ρ and Kendall's τ [32]. In our case the two ranked lists are represented by the *SLM* ranking r_1, and the user *consensus* ranking r_2. We discuss each of the correlation measures in more detail, where we will abbreviate *Spearman's footule* to

simply *Footrule* throughout the rest of this section.

The *Footrule* is calculated by summing the result of the absolute differences between the rank positions of the documents for each individual ranked list. The *Footrule*, denoted by Fr, is more formally defined as follows:

$$Fr(r_1, r_2) = \sum_{i=1}^{k} |(r_1(i) - r_2(i))| \tag{10}$$

where r_1 and r_2 are two ranked lists assumed to contain the same set of documents, and k is the size of the ranked list, in our case $k = 10$, which represents the top-10 ranked documents. In order to use the *Footrule* as a metric, we need to normalise the result by calculating the maximum possible value, through:

$$F = 1 - \frac{Fr(r_1, r_2)}{maxFr(k)} \tag{11}$$

where $maxFr$ represents the maximum value, which when k is an even number $maxFr = \frac{1}{2}k^2$, and if k is an odd number then $maxFr = \frac{1}{2}(k+1)(k-1)$. This ensures the resulting *Footrule* falls in the range of 0 and 1 where a value close to 1 means that the two ranked lists are highly similar.

When evaluating search results, however, we may wish to consider the fact that documents in the top ranks are often considered the most relevant to the users information need than documents appearing in lower ranks [21]. To give more weight to the top ranked documents, we apply the *M*-measure, which was designed to place more emphasis on ranked lists containing identical or near-identical sets of documents in the top rank positions. Due to the fact that the ranked lists contain the same set of documents, we can drop the terms S and T mentioned in [21], which record the set of documents unique to r_1 and r_2, respectively, and reformulate the *M*-measure more precisely as:

$$m = \sum_{i=1}^{k} |\frac{1}{r_1(i)} - \frac{1}{r_2(i)}|, \tag{12}$$

where we calculate the sum of the absolute differences between each document's *SLM* rank and *consensus* rank. Next we calculate the maximum value $max\ M$, which is defined through:

$$maxM = \sum_{i=1}^{k} |\frac{1}{i} - \frac{1}{k-i+1}| \tag{13}$$

Lastly, we normalise m by deducting 1 from the result of the division of m by the maximum value $maxM$ to obtain a metric ranging between 0 and 1:

$$M = 1 - \frac{m}{maxM(k)} \tag{14}$$

The average by query is calculated by summing the scores for each correlation measure, *Footrule* and *M*-measure, and then dividing the result by the total number of queries. We repeat this process for each user by first gathering the per query correlation scores for each user and then dividing the result by the total number of queries, and then take a further average for each user. The average *user agreement* represents the degree of correlation between each user and the *consensus* ranking, in other words, the *user agreement* describes the average correlation between an individual user and what we could consider to be the "wisdom" or "opinion" of the crowd. We produce a *consensus* ranking for each query, and measure the correlation between this ranking and the individual user ranking for the given query. Next, we calculate the average correlation per query for each user, as we did before, and then compute the average consensus based on the per user averages.

8.4.2 Normalised Discounted Cumulative Gain (NDCG)

While the correlation measures tell us how well the system generated ranking correlates with the user judgements, it does not directly describe the quality of the ranking algorithm. The ranking performance of an information retrieval system can be measured with the *Normalised Discounted Cumulative Gain* (*NDCG*) computed over the graded relevance scores. The $NDCG$ for a ranked list of size k is calculated through:

$$NDCG_k = \frac{DCG_k}{IDCG_k} \tag{15}$$

where DCG is the discounted cumulative gain, $IDCG$ represents the ideal DCG, obtained by sorting the documents in descending order by relevance value, and then calculating the DCG to get the maximum DCG, which is used in the normalisation step.

We selected two discounting functions for comparison. The first method reduces the contribution of the relevance score according to rank position, which we define as n. The second discounting function is the more common approach where the relevance score is discounted by $\log_2$ of the rank position. The DCG at a particular rank position k is defined as:

$$DCG_k(r) = rel_1 + \sum_{i=2}^{k} \frac{rel_i}{\log_2 i} \tag{16}$$

where r is a ranked list containing documents, and rel_i is the relevance score at position i and $\log_2 i$ represents the discounting function. The discounting function models user persistence [40] in terms of whether the user will continue to look for more documents further down the search results. This is achieved by reducing the contribution of the relevance score assigned to each document as a function of its position in the ranked list. Documents appearing later in the

ranked list are unlikely to be as relevant to the user as those in the top ranks and therefore only a small part of the document relevance score is passed on to the cummulative gain. The resulting $NDCG$ score ranges from 0 to 1, where a value of 1 means the ranking quality of the system is perfect as it is equivalent to the $IDCG$. We take the query results for each user and calculate the $NDCG$ and an average for each user, before calculating a final average over all queries.

8.4.3 Significance testing

As part of the assessment we evaluate the statistical significance of the results. We adopt the *bootstrap* method [29, 62, 59], which attempts to approximate the original underlying distribution of the population, by selecting a series of random samples of size N with replacement from the observed population data. An advantage of the bootstrap method is that it is compatible with any statistical measure [62], meaning we can use the correlation and $NDCG$ scores as our test statistics. Under the bootstrap method, we assume that the null hypothesis is that there is no difference between the ranking generated by the system and the ranking generated by the user evaluations. The difference is considered significant, with respect to the stated significance level, if the confidence intervals do not overlap. In order to obtain the confidence intervals we generate a series of samples by selecting a value at random from the original measures (*Footrule*, *M*-measure, and $NDCG$) to generate a sample equivalent in size to the number of queries or users in the original evaluation. The sampling process can be thought of as extracting values from the rows and columns of a n by m matrix, where the rows contain the correlation or $NDCG$ scores by query, and the columns represent the per user scores. Each random sample b, where $b = 1, ..., B$, is composed of values selected with replacement. We perform this operation for a total sample size of $B = 1000$ and calculate the average of the test statistic for each sample. Calculating the final confidence interval then involves sorting the averages in ascending order, and selecting the values that fall at the $B(1-(\alpha/2))$ and $B(\alpha/2)$ percentile, where α is the required significance level and $\alpha = 0.05$ represents a 95% confidence interval. We take an average over the lower and upper bounds of the confidence intervals and partition the results by query, user, and *user agreement* (see Section 8.5).

8.5 Evaluation Results

The evaluation was run over several days and 65 participants attempted the evaluation. A total of 24 users successfully completed the survey and the majority of the submissions were received from men between the ages of 20-30 years, and resident or born in North America. Of the toal submissions received, we excluded 10 users due to incomplete results caused by connection timeout issues, and 31 users who failed to pass the test queries, which is almost half of the total attempted submissions. It is interesting to note that 10 of the users who failed the test did not even pass the first test query. This means they were unable to identify that the top-10 results were composed of two completely dif-

ferent queries of different lengths (a short query versus a long verbose query), which highlights the importance of designing tests as part of an evaluation to filter potentially poor performing users.

8.5.1 Correlation measures

In this section we report the final correlation scores. Before continuing, we establish a baseline figure for each measure. We compute the *Footrule* and *M*-measure between the *SLM* ranking and the *display* order of the documents for each query, and then take the average over all queries to obtain an average baseline score. The baseline figures are reported in Table 1.

The baseline correlation for the 10 Samtla queries is 1.0 since the display and

Baseline *Correlation*		
Type	*Footrule*	*M-measure*
Random queries (40)	0.400	0.369

Table 1: Baseline values for each correlation measure divided by query type

the system ranking are equivalent and is not included here. In terms of the baseline for the 40 *random* order queries, we can see that it is quite low across the two measures, but does show that there is some correlation despite the shuffling process. The final results for each measure are displayed in Table 2 and Table 3 where we present the average *Footrule* and *M*-measure for the *SLM* and *display* ranking compared to the user *consensus* ranking, respectively. For each form of analysis, we divide the results into query, user, and user consensus averages, and report the 95% confidence interval in square brackets, obtained from the bootstrap (see Section 8.4.3).

	SLM	
Samtla queries (10)	*Footrule*	*M-measure*
Query	0.775 [0.775 - 0.779]	0.840 [0.840 - 0.844]
User	0.863 [0.859 - 0.863]	0.906 [0.903 - 0.906]
User consensus	0.800 [0.795 - 0.798]	0.847 [0.845 - 0.848]
Random queries (40)	*Footrule*	*M-measure*
Query	0.757 [0.756 - 0.759]	0.761 [0.759 - 0.763]
User	0.853 [0.851 - 0.854]	0.846 [0.844 - 0.847]
User consensus	0.716 [0.715 - 0.718]	0.737 [0.735 - 0.739]

Table 2: Average correlation scores for the *random* queries ordered by the *SLM* ranking divided into query, user, and user consensus

The results of Table 2 show that the user relevance judgments for the 10 Samtla queries is comparable or higher for the Footrule and the M-measure, implying that a display bias may be present. In particular, we see that the average score

by user for the M-measure is 0.906, suggesting that users were more likely to assign higher relevance grades to a selection of documents appearing at the top of the search results. Consequently, we discard this set of queries from the analysis, and focus solely on the *random* order queries for the remainder of this section.

The average correlation scores for the 40 *random* queries are positively correlated to the *consensus* ranking. Looking at the *Footrule*, we see that the query average of 0.757 is lower than the user average of 0.853. This may be attributed to query length, in terms of the fact that longer queries had an average of 3 to 4 highly relevant documents, with the remaining documents containing only partial matches. On the other hand, short keyword queries tended to have full matches to the query in all the documents retrieved and so users assigned proportionally higher relevance grades across more documents in the search results. We also observe similar results for the average M-measure, which shows that users were assigning more relevance to a selection of documents that are ranked in the top positions according the underlying SLM.

The average *user consensus* also suggests that the majority of users were in agreement with the crowd opinion of which documents were the most relevant. Turning to the correlation scores calculated over the *display* ranking, the results in Table 3 summarise the final averages for the *Footrule* and M-measure over the 40 *random* order queries according to the order in which the documents were presented to the users.

We can see that the correlation is much lower than both the 10 *Samtla* queries

Random queries (40)	*Footrule*	*M-measure*
Query	0.402 [0.400 - 0.404]	0.474 [0.471 - 0.476]
User	0.435 [0.434 - 0.437]	0.416 [0.415 - 0.417]
User consensus	0.660 [0.658 - 0.661]	0.669 [0.668 - 0.671]

Table 3: Average correlation scores for the *display* ranking divided into query, user, and average user consensus for the *random* order queries.

and the *random* order queries ranked according to the SLM (Table 2) across queries and users. This suggests that the users were less sensitive to the presentation order of the documents, in other words they were attempting to do a good job rather than assigning relevance as a function of the document position. The average correlation scores for the *display* order of the 40 *random* order queries are still positively correlated, suggesting that there is some presentation bias. However, if we take in to account the degree of correlation that already existed between the two permutations on the order of documents, that is 0.400 and 0.369 for the *Footrule* and M-measure, respectively (see Table 1), then we could argue that the average correlation scores are actually much smaller. Taking this issue in to account, we can conclude that there was not much bias in the users'

judgements in terms of the presentation order of the documents.

To summarise, there is an observable difference in the correlation scores for the *SLM* order and the *display* order, with the correlation measures for the *SLM* being higher than the *display* order suggesting that users agreed with the ranking generated by the *SLM*. In general, it appears that the crowd of users were not affected by the presentation order of the documents when they were in *random* order, and we are able to reject the null hypothesis that there is no correlation between the *SLM* system ranking and the ranking generated from the user relevance judgements, which is supported by the fact that the confidence intervals do not overlap meaning we have a significant result at the 95% confidence level.

8.5.2 Normalised Discounted Cumulative Gain (NDCG)

Before presenting the results for the $NDCG$ measure, we establish a baseline once again establish a baseline. This is done using a different method to the correlation measures, where we simulate the input provided by 1000 random users. Each user is represented by a random assignment of relevance grades to the documents for each query, which we then summarise by computing the average $NDCG$ by user and query for each discounting function n and log_2, respectively, which are presented in Table 4 below.

Baseline $NDCG$	
n	log_2
0.853	0.870

Table 4: Baseline NDCG

The average baseline figures are fairly close to the maximum $NDCG$, with the logarithmic discounting function log_2 being slightly less aggressive than the discounting by rank position n. We have found that we obtain similar results regardless of the adopted discounting function, however, we include both as they provide different models of user persistance, with n representing a more impatient user. The final average scores for the $NDCG$ applied to the *SLM* ranking and the *display* ranking are presented below (see Table 5 and Table 6). We make a distinction between the *display* ranking for the 10 *Samtla* queries and the 40 *random* order queries and report the total average by query and user with their 95% confidence intervals presented alongside in square brackets.

As with the correlation scores, we can see that the users tended to assign higher relevance to the top documents in the search results, as illustrated by the $NDCG$ scores being very close to the $IDCG$. The query and user averages based on discounting by n are equal as a result of rounding, but there was only a slight difference between the two (0.98573 and 0.98564 respectively). As mentioned,

NDCG@10	**SLM**	
Samtla queries (10)	n	log_2
Query	0.985 [0.985 - 0.985]	0.988 [0.987 - 0.988]
User	0.985 [0.985 - 0.985]	0.987 [0.987 - 0.987]
Random queries (40)	n	log_2
Query	0.981 [0.980 - 0.981]	0.983 [0.983 - 0.984]
User	0.982 [0.981 - 0.982]	0.984 [0.983 - 0.984]

Table 5: Average correlation scores for the *SLM* ranking divided in to query and user averages with the 95% confidence intervals in square brackets.

we can see that the $NDCG$ is slightly higher for the 10 Samtla queries than the 40 random queries, suggesting that users assigned higher scores to the top documents as a result of their position being at the top of the search results. Consequently, we remove these 10 queries from the results and discussion, as there would once again appear to be a slight presentation bias.

Turning to the average $NDCG$ scores for the *SLM* ranking of the 40 *random* order queries, we see that the average query and user $NDCG$ scores are quite close for the *SLM* ranking (0.983 by query, and 0.984 by user assuming a discounting function of log_2). If we compare these results to the $NDCG$ scores of *display* order of the queries (see Table 6) and the baseline $NDCG$ scores (see Table 4), there was less of a presentation bias due to the relatively low $NDCG$ at 0.881 for the discounting function by rank position n, and 0.894 for the discounting function of log_2. This means that users were not heavily influenced by the presentation order of the documents. We observe that the 95% confidence intervals from the *bootstrap* process do not overlap, which means we can conclude that the results are significant at the $\alpha = 0.05$ level. And we noted that the users gave more relevance to documents appearing in the top ranks of the *random* order queries, represented by the high average $NDCG$ score for the *SLM* ranking (Table 6). Naturally, these results assume that we take in to account the baseline $NDCG$ scores, which means the $NDCG$ for the *display* order is actually much lower.

To summarise, on the basis of the performance measures presented above,

NDCG@10	**Display**	
Random queries (40)	n	log_2
Query	0.881 [0.881 - 0.883]	0.894 [0.894 - 0.896]
User	0.882 [0.881 - 0.883]	0.895 [0.894 - 0.896]

Table 6: Average correlation scores for the*display* ranking divided in to query and user averages for the 40 *random* order queries, with the 95% confidence intervals in square brackets.

we can say that users were highly correlated with the *SLM* order of the queries than the *display* order when analysing the results of the *random* order queries independently of the 10 *Samtla* queries. The users were more influenced by the presentation order of the 10 *Samtla* queries, in the sense that they were slightly more generous with their relevance grades, where they tended to assign higher relevance to a few documents at the very top of the search results shown by the high M-measure and $NDCG$ scores. Out of the 50 queries completed by each user, 80% of them were presented in random order, yet we see that the users consistently assigned more relevance to the documents that received the highest document score according to the underlying SLM, and we can see that these scores are not the result of users assigning relevance at random, or "gaming" the system, in part due to the role played by the quality assessment represented by the test queries. We also observed users revisiting their earlier relevance assignments, when they encountered highly relevant documents at the bottom of the result page, caused by the random shuffle process. Therefore, there is significant evidence to suggest that users were attempting to do a good job and were not assigning relevance grades purely at random, but based on what they considered to be relevant given the provided query context.

We can conclude then, that the ranking quality of *Samtla* and its underlying SLM correlates well with the ranking generated by the user relevance judgements, both in terms of which documents were relevant and also of the top document, which were most likely to meet their information across query types, from single word queries to long more verbose queries.

8.6 Discussion

Crowdsourcing has its challenges, in particular, the researcher has little control over the evaluation process once it is launched and available online. Therefore, as we have demonstrated, it is necessary to consider the use of test queries in order to filter out bad users upfront e.g. those who have not understood the task or do not have the correct attitude. This increases the quality of the submissions, and mitigates against issues that can arise, such as an unhappy user as a result of a rejected submission, or withholding payment due to a suspect submission. These issues can be difficult to resolve and may have an impact on your reputation, and consequently on whether you will be able to submit future evaluations with the same crowd sourcing platform.

The design of the evaluation should record data that permits the testing of a display bias, since some users may assign relevance to document in the top ranks without necessarily digesting the snippets fully. This is easily achievable by randomising the order of the queries. It is also worth recording a timestamp for each response. This enables the researcher to check for users who are speeding through the evaluation at a rate that exceeds the ability to comfortably digest the information related to the task. We found that users assigned relevance at an average rate of three seconds per rank position. The minimum

time taken was one second, which we could argue is not enough time to digest the snippet and then navigate to the drop-down box to select a relevance grade. The maximum time to select a relevance grade was 13 minutes, but this is likely the result of users being interrupted or distracted from the task.

Furthermore, the difference between the total query average and user averages can be explained by the fact that users tended to adopt their own strategy for assigning relevance. A large number of users did not make use of all relevance grades (see Figure 14), but instead adopted a binary relevance approach where they only assigned grades of "Very Relevant" or "Not Relevant" to the documents. The short queries tended to have more relevant documents in the top-10 meaning that the user tended to judge relevance based on the total number of highlighted terms in the snippet. On the other hand, the longer verbose queries contained an average of 3 to 4 "Very Relevant" documents, with the remaining results containing partial matches to the query, which received less relevance. For example, documents containing a full match for the query "...[As the Lord commanded]..." naturally received higher relevance scores than the partial match "...[As th]y [Lord commanded]...". However, this is often user-dependent, and it could be argued that a researcher of the Bible would find the latter example just as relevant to their information need, or at least, that it provides an interesting example for their research.

In conclusion, we have shown that non-parametric correlation and $NDCG$ measures provide a good basis for assessing the performance of an information retrieval system. The non-parametric correlation measures show the degree of agreement between what users considered relevant and the ranking generated by the SLM (see Section 4). On the other hand, the $NDCG$ described the ranking quality of the ranked lists, and we observe that the system consistently produces a ranking where the top ranks are occupied by the most relevant documents. We also described how we can measure the overall opinion or agreement between the users by comparing each user with the *consensus* ranking, which showed that each individual user agreed on average with the ranking generated by the crowd. Lastly, the significance of the results was evaluated with the *bootstrap* method, which is non-parametric, relatively simple to implement, and as effective as other significance tests [59].

Using crowdsourcing as a platform for system evaluation provides researchers with access to a large group of potential participants, but as we have demonstrated, it is necessary to design the evaluation in such a way so as to minimise technical challenges, minimise poor quality results, and record data on user interaction with the evaluation software in order to spot potential cheating.

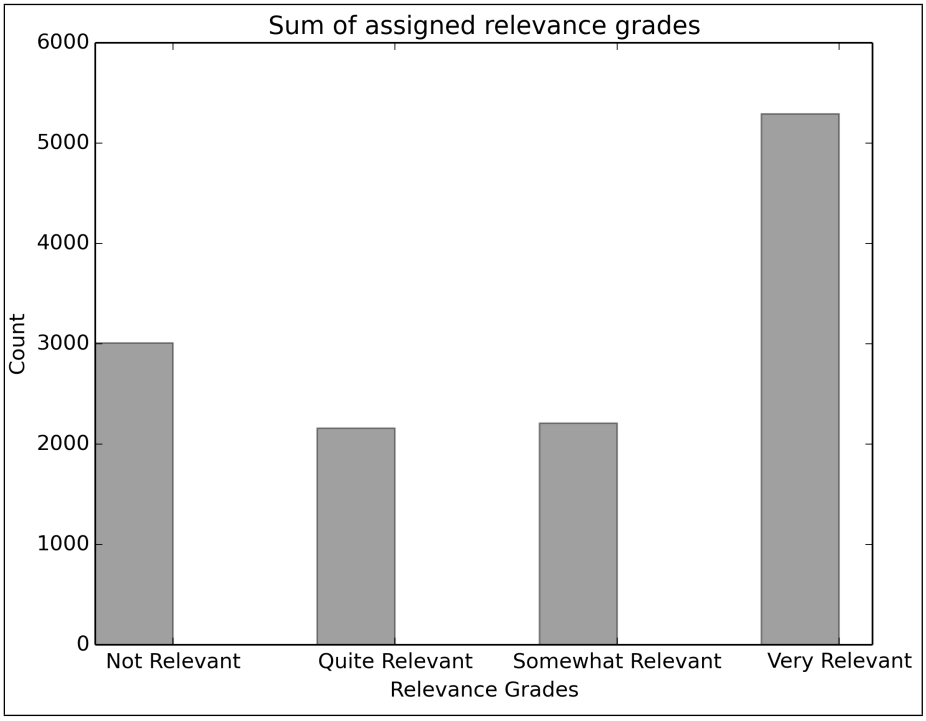

Figure 14: Distribution of relevance grades used in the evaluation.

9 Concluding Remarks

We have introduced the underlying framework of the Samtla system (Section 3), and the data structures and algorithms adopted. We showed how statistical language models can be used for ranking documents according to user queries (Section 4) and demonstrated that our implementation is providing users with the most relevant documents in the top ranks of the search results (Section 8).

We also described how users interact with the system (Section 6), and the tools we have currently released to our user groups (Section 5). The case studies provide an insight into how our users are currently using these tools to carry out their research (Section 7).

We are now focusing on the development of the underlying framework where we look at additional parameters that can be incorporated in to the data model (see Section 4) in order to add a layer of semantics to the search component. For example, we currently assume a uniform prior for all document probabilities when ranking the documents in response to a query. We can use the *JSD* matrix generated for the related documents tool (see Section 5.2) to compute a non-uniform prior, which will enable us to integrate document-specific knowledge as

part of the Samtla query model. Further work is centered on simple methods for identifying important events in the collection documents, which could be presented to users as a timeline. This task is often referred to as event tracking and identification [60].

The main novelty of Samtla is the underlying probabilistic model that has enabled us to develop a diverse range of tools that are language independent and applicable to many document collections, including flexible search and mining of text patterns, document comparison, query and document recommendation, and the way the system can incorporate external sources of information in the form of metadata provided by users or third-party sources such as Wikipedia, to supplement the toolset. Samtla aims to complement existing methods in the digital humanities by helping researchers with their research needs by providing a general purpose environment.

In summary, we have discussed how systems developed for the Humanities can be made 'future-proof' in the sense of providing a generalised framework that can be easily extended to new document collections without changes to the underlying system components in order to compensate for language-specific issues such as word stemming and tokenisation. Although Samtla is still in development we already have a number of Samtla systems available for a range of document collections (King James Bible, Aramaic Magic Bowls, Vasari, the Microsoft Corpus, and the Financial Times), which cover a broad range of corpora composed of one or more languages including Aramaic, Syriac, Mandaic, Hebrew, English, German, French, Hungarian, Italian, and Russian. In addition, Samtla is not necessarily restricted to historic document collections, but can be extended straightforwardly to other application domains, which require search and mining of text patterns, such as medical and legal text collections.

References

[1] Cultivating understanding through research and adaptivity. [Online; accessed 20 February 2015].

[2] Ibm languageware. [Online; accessed 18 February 2015].

[3] What is the 1641 depositions project? [Online; accessed 20 February 2015].

[4] What is flat design? `http://gizmodo.com/what-is-flat-design-508963228`, 2013. [Online; accessed 31-January-2014].

[5] The history of flat design: How efficiency and minimalism turned the digital world flat. `http://thenextweb.com/dd/2014/03/19/history-flat-design-efficiency-minimalism-made-digital-world-flat/`, 2014. [Online; accessed 17-May-2015].

[6] Prolific academic. https://prolificacademic.co.uk/, 2014. [Online; accessed 28-January-2014].

[7] Vmba: Virtual magic bowl archive. http://www.southampton.ac.uk/vmba/, 2014. [Online; accessed 28-January-2014].

[8] Accordance. http://www.accordancebible.com, 2015. [Online: accessed 01-October-2015].

[9] The django project. https://www.djangoproject.com/, 2015. [Online; accessed 12-January-2015].

[10] Ecma-404 the json data interchange standard. http://json.org/, 2015. [Online; accessed 12-January-2015].

[11] Html5 - a vocabulary and associated apis for html and xhtml. http://www.w3.org/TR/html5/, 2015. [Online; accessed 12-January-2015].

[12] jquery. write less do more. http://jquery.com/, 2015. [Online; accessed 12-January-2015].

[13] List of biblical places. https://en.wikipedia.org/wiki/List_of_biblical_places, 2015. [Online; accessed 02-May-2015].

[14] List of major biblical figures. https://en.wikipedia.org/wiki/List_of_major_biblical_figures, 2015. [Online; accessed 02-May-2015].

[15] Lists of lists on wikipedia. https://en.m.wikipedia.org/wiki/List_of_lists_of_lists, 2015. [Online; accessed 01-October-2015].

[16] Responsa - features. http://www.biu.ac.il/JH/Responsa/features.htm, 2015. [Online; accessed 27-October-2015].

[17] Sorensen, lee. "vasari, giorgio." in dictionary of art historians. http://www.dictionaryofarthistorians.org/wittkowerr.htm, Retrieved 28th June 2015.

[18] C. Aggarwal and C. Zhai. *Mining Text Data*. Springer-Verlag New York Inc, 2012.

[19] O. Alonso and R. A. Baeza-Yates. Design and implementation of relevance assessments using crowdsourcing. In *Advances in Information Retrieval - 33rd European Conference on IR Research, ECIR 2011, Dublin, Ireland, April 18-21, 2011. Proceedings*, pages 153–164, 2011.

[20] J. Bar-Ilan, K. Keenoy, M. Levene, and E. Yaari. Presentation bias is significant in determining user preference for search resultsa user study. *Journal of the American Society for Information Science and Technology*, 60(1):135–149, 2009.

[21] J. Bar-Ilan, M. Mat-Hassan, and M. Levene. Methods for comparing rankings of search engine results. *Comput. Netw.*, 50(10):1448–1463, July 2006.

[22] D. M. Bikel, R. Schwartz, and R. M. Weischedel. An algorithm that learns what's in a name. *Mach. Learn.*, 34(1-3):211–231, Feb. 1999.

[23] D. C. Brabham. Crowdsourcing as a model for problem solving: An introduction and cases. *Convergence*, 14(1):75, 2008.

[24] V. Broughton. Faceted classification as a basis for knowledge organization in a digital environment: The bliss bibliographic classification as a model for vocabulary management and the creation of multidimensional knowledge structures. *New Rev. Hypermedia Multimedia*, 7(1):67–102, July 2002.

[25] M. Bruls, K. Huizing, and J. van Wijk. Squarified treemaps. In *In Proceedings of the Joint Eurographics and IEEE TCVG Symposium on Visualization*, pages 33–42. Press, 1999.

[26] S. Chen and J. Goodman. An empirical study of smoothing techniques for language modeling. *Computer Speech and Language*, 13(4):359–394, 1999.

[27] Y. Choueka. Responsa: A full-text retrieval system with linguistic processing for a 65-million word corpus of jewish heritage in hebrew. *Data Eng.*, 12(4):22–31, Nov. 1989.

[28] S. P. Crain, K. Zhou, S.-H. Yang, and H. Zha. Dimensionality reduction and topic modeling: From latent semantic indexing to latent dirichlet allocation and beyond. In C. C. Aggarwal and C. Zhai, editors, *Mining Text Data*, pages 129–161. Springer, 2012.

[29] A. C. Davison and D. V. Hinkley. *Bootstrap Methods and their Application*. Cambridge University Press, 1997.

[30] P. Diaconis and R. L. Graham. Spearman's footrule as a measure of disarray. *Royal Statistical Society Series B*, 32(24):262–268, 1977.

[31] D. M. Endres and J. E. Schindelin. A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860, 2003.

[32] R. Fagin, R. Kumar, and D. Sivakumar. Comparing top k lists. In *In Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 28–36, 2003.

[33] U. M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. Advances in knowledge discovery and data mining. chapter From Data Mining to Knowledge Discovery: An Overview, pages 1–34. American Association for Artificial Intelligence, Menlo Park, CA, USA, 1996.

[34] H. Garcia-Molina, M. Joglekar, A. Marcus, A. Parameswaran, and V. Verroios. Challenges in data crowdsourcing. *Knowledge and Data Engineering, IEEE Transactions on*, PP(99):1–1, 2016.

[35] F. Gibbs and T. Owens. Building better digital humanities tools: Toward broader audiences and user-centered designs. *Digital Humanities Quarterly*, 6(2), 2012.

[36] D. Gusfield. *Algorithms on Strings, Trees, and Sequences: Computer Science and Computational Biology*. Cambridge University Press, 1997.

[37] M. A. Hearst. Uis for faceted navigation: Recent advances and remaining open problems. In *in the Workshop on Computer Interaction and Information Retrieval, HCIR 2008*, 2008.

[38] M. A. Hearst. *Search User Interfaces*. Cambridge University Press, 1 edition, 2009.

[39] S. Huston and W. B. Croft. Evaluating verbose query processing techniques. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '10, pages 291–298, New York, NY, USA, 2010. ACM.

[40] K. Järvelin and J. Kekäläinen. Cumulated gain-based evaluation of ir techniques. *ACM Trans. Inf. Syst.*, 20(4):422–446, Oct. 2002.

[41] J. Kazama and K. Torisawa. Exploiting wikipedia as external knowledge for named entity recognition. In *Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 698–707, 2007.

[42] A. Kittur, E. H. Chi, and B. Suh. Crowdsourcing user studies with mechanical turk. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 453–456, New York, NY, USA, 2008. ACM.

[43] V. Lavrenko, M. D. Schmill, D. Lawrie, P. Ogilvie, D. Jensen, and J. Allan. Language models for financial news recommendation. In *CIKM*, pages 389–396. ACM, 2000.

[44] M. Lease and E. Yilmaz. Crowdsourcing for information retrieval: introduction to the special issue. *Information Retrieval*, 16(2):91–100, 2013.

[45] A. Leff and J. Rayfield. Web-application development using the model/view/controller design pattern. In *Enterprise Distributed Object Computing Conference, 2001. EDOC '01. Proceedings. Fifth IEEE International*, pages 118–127, 2001.

[46] M. Levene. *An Introduction to Search Engines and Web Navigation*. John Wiley & Sons, Hoboken, New Jersey, 2nd edition, 2010.

[47] D. D. Lewis. Naive (bayes) at forty: The independence assumption in information retrieval. In *Proceedings of the 10th European Conference on Machine Learning*, ECML '98, pages 4–15, London, UK, UK, 1998. Springer-Verlag.

[48] J. Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37:145–151, 1991.

[49] X. Liu and W. B. Croft. Statistical language modeling for information retrieval. *Annual Review of Information Science and Technology*, 39(1):1–31, 2005.

[50] J. Ma and L. Zhang. Modern BLAST programs. In *Problem Solving Handbook in Computational Biology and Bioinformatics*. Springer US, 2011.

[51] C. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.

[52] P. Mcnamee and J. Mayfield. Character n-gram tokenization for european language text retrieval. *Inf. Retr.*, 7(1-2):73–97, Jan. 2004.

[53] N. Megiddo and D. S. Modha. Outperforming lru with an adaptive replacement cache algorithm. *IEEE Computer*, 37(4):58–65, 2004.

[54] S. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge University Press, New York, NY, USA, 2nd edition, 2009.

[55] S. Mizzaro. Relevance: The whole history. *JOURNAL OF THE AMERICAN SOCIETY FOR INFORMATION SCIENCE*, 48:810–832, 1997.

[56] D. Nadeau and S. Sekine. A survey of named entity recognition and classification. *Lingvisticae Investigationes*, 30(1):3–26, 2007.

[57] T. Rommel. Literary studies. In *A Companion to Digital Humanities*, pages 88–96. Blackwell Publishing Ltd, 2007.

[58] R. Rosenfeld. Two decades of statistical language modeling: Where do we go from here? In *Proceedings of the IEEE*, volume 88, pages 1270–1278, 2000.

[59] T. Sakai. Evaluating evaluation metrics based on the bootstrap. In *Proceedings of the 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '06, pages 525–532, New York, NY, USA, 2006. ACM.

[60] H. Sayyadi, M. Hurst, and A. Maykov. Event detection and tracking in social streams. In *In Proceedings of the International Conference on Weblogs and Social Media (ICWSM 2009). AAAI*, 2009.

[61] M. H. Schulz, S. Bauer, and P. N. Robinson. The generalised k-truncated suffix tree for time-and space-efficient searches in multiple dna or protein sequences. *IJBRA*, 4(1):81–95, 2008.

[62] M. D. Smucker, J. Allan, and B. Carterette. A comparison of statistical significance tests for information retrieval evaluation. In *Proceedings of the Sixteenth ACM Conference on Conference on Information and Knowledge Management*, CIKM '07, pages 623–632, New York, NY, USA, 2007. ACM.

[63] M. Steinbach, G. Karypis, and V. Kumar. A comparison of document clustering techniques. In *In KDD Workshop on Text Mining*, 2000.

[64] M. Sweetnam, M. Agosti, N. Orio, C. Ponchia, C. Steiner, E. Hillemann, M. Ó Siochrú, and S. Lawless. User needs for enhanced engagement with cultural heritage collections. In *Proceedings of the International Conference on Theory and Practice of Digital Libraries (TPDL)*, pages 64–75, 2012.

[65] M. L. Wilson. Search user interface design. *Synthesis Lectures on Information Concepts, Retrieval, and Services*, 3(3):1–143, 2011.

[66] O. F. Zaidan and C. Callison-Burch. Crowdsourcing translation: Professional quality from non-professionals. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies - Volume 1*, HLT '11, pages 1220–1229, Stroudsburg, PA, USA, 2011. Association for Computational Linguistics.

[67] C. Zhai. *Statistical Language Models for Information Retrieval*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers, San Francisco, 2009.

[68] C. Zhai and J. Lafferty. The dual role of smoothing in the language modeling approach. In *Proceedings of the Workshop on Language Models for Information Retrieval (LMIR) 2001*, pages 31–36, 2001.

[69] C. Zhai and J. Lafferty. A study of smoothing methods for language models applied to information retrieval. *ACM Trans. Inf. Syst.*, 22(2):179–214, Apr. 2004.